Growth Lab

Japan's Innovation Challenge: Escaping the Middle-Technology Trap

Dany Bahar, Shreyas Gadgin Matha, Ricardo Hausmann, and Santiago Segovia

Growth Lab Working Paper Series
No. 269

May
2026

GROWTH LAB
HARVARD KENNEDY SCHOOL
79 JFK STREET
CAMBRIDGE, MA 02138

GROWTHLAB.HKS.HARVARD.EDU

This report may be referenced as follows: Bahar, D., Gadgin Matha, S., Hausmann, R., Segovia, S. (2026). “Japan’s Innovation Challenge: Escaping the Middle-Technology Trap.” Growth Lab Working Paper, John F. Kennedy School of Government, Harvard University.

Japan's Innovation Challenge: Escaping the Middle-Technology Trap

Dany Bahar* Shreyas Gadgin Matha† Ricardo Hausmann‡
Santiago Segovia§

May 2026

Abstract

Japan remains one of the world's most technologically sophisticated economies, yet its labor productivity has been stagnant for more than two decades. This paper investigates the apparent contradiction between Japan's high R&D intensity and its weak productivity performance by examining the allocation, composition, and effectiveness of innovation across industries. Using industry-level data from the OECD, patent-level data linked across technology and industry classifications, and a set of nine technological taxonomies, we document that Japan disproportionately concentrates R&D in mid-technology manufacturing sectors—such as motor vehicles, electrical equipment, and chemicals—that generate relatively low productivity spillovers. High-technology sectors, including ICT, pharmaceuticals, scientific R&D, and advanced digital services, receive a significantly smaller share of investment and exhibit much higher productivity contributions in other countries. We further show that Japan's indirect, tax-based system of R&D support reinforces this equilibrium by favoring large incumbents and under-supporting SMEs. We conclude by assessing the potential of Japan's new 17-sector strategy to reorient the innovation system toward frontier technologies.

*Center for Global Development; Brown University; and Harvard Growth Lab
†Harvard Growth Lab
‡Harvard University; Harvard Growth Lab; and CEPR
§Harvard Growth Lab

Introduction

Japan occupies an unusual position in the global economy. It is a highly advanced, innovative country with a deep stock of productive knowledge in the world, a central role in global value chains, and an industrial structure that remains sophisticated by international comparison. Yet over the past two decades, aggregate labor productivity growth has been weak. Between 2000 and 2021, Japan's labor productivity growth was essentially flat on average, while the United States, Germany, and other high-income economies recorded steady annual gains. The gap in productivity levels has widened over time, leaving Japan richer in its knowledge base than many peers, but poorer in productivity growth.

This divergence motivates the analysis in this paper. Japan is not a low-innovation economy: it remains near the technological frontier, sustains world-leading R&D expenditure—over 2.5% of GVA—and maintains export strengths in engineering, automobiles, robotics, and precision manufacturing. However, output per worker has been largely stagnant for twenty years. Understanding this disconnect between innovation effort and productivity performance is central to explaining Japan's growth record.

Our focus is on the allocation of innovative activity across sectors. Productivity growth at the frontier requires not only investment in knowledge, but effective allocation of that investment toward sectors capable of generating new value at scale. This is particularly important in economies like Japan that have largely exhausted gains from technology adoption and catch-up: when imitation margins are limited, growth depends more on generating and scaling *new* technologies than on incremental improvements to existing ones. In Schumpeterian models of growth through creative destruction, growth is driven by quality-improving innovations that render existing technologies obsolete, and frontier performance depends on the arrival rate and size of such innovations rather than on the mere intensity of R&D spending.[1] This perspective highlights that what matters for productivity is not just how much a country spends on R&D, but where in the economy that spending occurs.

We ask whether Japan's innovative effort is concentrated in sectors with relatively modest productivity payoffs. Using data on private R&D over two decades, we document that Japan allocates an unusually large share of R&D to mid-technology industries. More than half of private R&D takes place in sectors such as automobiles, electrical equipment, and industrial machinery—industries that are technologically intensive but structurally mature. High-technology sectors receive only 35–40% of private R&D, a substantially

[1] See, for example, Aghion and Howitt's model of growth through creative destruction, in which step-by-step innovations on a quality ladder generate long-run growth by replacing older technologies (Aghion and Howitt, 1992), and their subsequent synthesis of Schumpeterian growth theory (Aghion and Howitt, 2009).

lower share than in comparable economies such as the United States. This pattern persists across a wide range of alternative technology classifications, including industry-based taxonomies and patent-based measures.

This composition matters because mid-technology sectors tend to generate more incremental than transformative innovations: they optimize and refine existing production systems, but less frequently generate large productivity spillovers. By contrast, high-technology sectors—including ICT, AI-related activities, semiconductors, biotechnology, and advanced digital services—are more likely to produce general-purpose technologies and innovations with broad spillovers across the economy. In the Schumpeterian literature, frontier growth is associated with the process of creative destruction, in which new technologies displace old ones and reallocate resources toward more productive activities (e.g., Aghion and Howitt, 1992, 2009). From this perspective, an innovation system that directs a relatively small share of effort toward such breakthrough industries may struggle to sustain productivity growth even if aggregate R&D intensity is high.

Within this framework, our contribution is threefold. First, we document the structure of Japan's private R&D allocation over the past two decades and benchmark it against other major economies. Using a consensus classification that aggregates several technology taxonomies, we show that Japan is unusually concentrated in mid-technology manufacturing industries and comparatively underexposed to high-technology sectors. Second, we quantify how high-, mid-, and low-technology industries contribute to aggregate productivity growth across countries. Our regression analysis, implemented under nine alternative classification schemes, consistently finds that high-technology industries are associated with the largest productivity contributions, whereas mid-technology industries contribute significantly less on average. We then use this framework to assess whether Japan's emphasis on mid-tech industries reflects efficient specialization or misallocation, by comparing the productivity returns to R&D and the relative innovation efficiency of Japanese industries with those in peer economies. Finally, we examine the role of public policy in shaping this allocation. In particular, we argue that Japan's reliance on indirect, tax-based R&D incentives—which largely benefit large incumbents—tends to reinforce the existing mid-technology equilibrium, and we outline policy options for rebalancing innovation toward frontier sectors and high-tech, early-stage firms without undermining Japan's established industrial strengths.

Our main results show that mid-technology sectors in Japan account for a larger share of R&D than would be implied by their contribution to productivity growth, while high-technology sectors—though smaller in scale—exhibit higher returns per unit of R&D and greater innovation efficiency. The implication is not that Japan should reduce support for its industrial backbone, but that the marginal yen of additional R&D is likely to yield higher productivity gains if directed toward high-technology and younger, frontier-oriented firms.

The remainder of the paper is organized as follows. Section 2 describes the data, the construction of our industry panel, and the technology classifications used in the analysis. Section 3 presents stylized facts on Japan's productivity performance in comparative perspective. Section 4 turns to the structure of Japan's innovation system, documenting the allocation of R&D across sectors and technology levels and assessing the robustness of our findings across alternative taxonomies. Section 5 evaluates whether Japan's pattern of specialization is a feature or a misallocation by relating sectoral R&D allocation to productivity contributions and innovation efficiency, and examining the role of public policy in sustaining the current equilibrium. Section 6 analyzes the Japanese government's new 17–sector strategy as a potential vehicle for reallocating innovation effort toward frontier technologies. Section 7 concludes with policy implications for rebalancing innovation, strengthening productivity, and enabling Japan to escape what Fuest et al. (2024) term the middle-technology trap.

Data Section

The primary data source for our analysis is the Organization for Economic Co-operation and Development (OECD). Specifically, we use data on gross value added (GVA), employment, and private research and development (R&D) expenditures by economic activity, as classified under the fourth revision of the International Standard Industrial Classification of All Economic Activities (ISIC Rev. 4). We further supplement this with patent data from the European Patent Office's PATSTAT database, processed by Chacua (2019) and Chacua (2023). We consider patent family count aggregations by inventor country, year, and technology classification at the four-digit International Patent Classification (IPC4) subclass level. We employ fractional counting, which weights each patent family by the inverse of the product of unique inventor countries and unique IPC4 classes assigned to the family (Miguélez et al., 2019). This approach ensures that multi-country and multi-IPC code patents are appropriately distributed across locations and technology domains, with the sum of fractional counts equaling the total number of patent families. With these sources, we calculate labor productivity—ratio between GVA and employment—at the industry level, as well as other innovation metrics (*e.g.*, R&D intensity, patenting efficiency, among others).

ISIC is the international reference classification of economic activities compiled by the United Nations. These economic activities are subdivided into a hierarchical, four-level structure of mutually exclusive categories. The categories at the highest level are called sections, which are alphabetically coded categories. The sections subdivide the entire spectrum of productive activities into broad groupings, such as "Agriculture, forestry and fishing" (section A), "Manufacturing" (section C), and "Information and communication" (section J). The classification is then organized into more detailed categories, which are numerically coded: two-digit divisions, three-digit groups, and, at the greatest level of detail, four-digit classes. The most granular level of data available for our analysis is at the division level. Nevertheless, not all countries report data to the OECD with the same level of detail or consistency.[2] In order to make international comparisons and maintain good data coverage across time, we create a balanced panel using a mix of section- and division-level data. Our balanced panel has approximately 19,000 observations, consisting of 24 countries[3] (mostly OECD member countries) and 36 industries[4] (7 sections and 29 divisions) between 2000 and 2021.

To analyze Japan's innovation specialization patterns, we employ patent data in two complementary ways. First, we classify patents by technological complexity to assess whether Japan exhibits a middle-technology trap in its innovation output. Since patent data are classified by technology (IPC) while industry R&D data are classified by economic activity (ISIC), we construct a bridge between these classification systems using empirical co-occurrence data. Specifically, we link IPC technology codes at the 4-digit level to NAICS industry codes at the 6-digit level using co-occurrences derived from the DISCERN 2.0 dataset (Arora et al., 2024), which links U.S. publicly listed firms to their patents and provides firm-level industry classifications (*i.e.*, NAICS codes). We use this firm-level linkage to capture which industries produce patents in specific technology classes, reflecting organizational capabilities that span industrial and technological domains. We then classify each NAICS industry as high-tech, mid-tech, or low-tech based on Product Complexity Index (PCI) values calculated from U.S. County Business Patterns data, following the methodology of Hausmann et al. (2014). Specifically, we classify NAICS industries at or below the 25th percentile of PCI as low-tech, those at or above the 75th percentile as high-tech, and those in between as mid-tech. Using the IPC-NAICS co-occurrence matrix, we propagate industry complexity to patent classes by computing each IPC4 code's weighted average PCI based on the industries that produce patents in that technology domain. We then classify each IPC4 code into complexity tiers using the same percentile thresholds (25th and 75th percentiles) applied to the distribution of IPC-level weighted average PCI values. For our patent-based analyses, we focus on the 2018–2020 period to ensure recent data with good coverage.

[2]For Japan specifically, we improved data coverage by leveraging data from RIETI's Japan Industrial Productivity (JIP) Database. The JIP database uses its own industrial classification with 100 divisions. We used a concordance table between JIP and ISIC classifications created by RIETI, which allowed us to increase the division-level data available for Japan from 12 divisions in the OECD data to 56.

[3]The countries in our sample are the following: Austria, Belgium, Bulgaria, Croatia, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Italy, Japan, Latvia, Mexico, Netherlands, Portugal, Romania, Slovak Republic, Slovenia, Spain, United Kingdom, United States

[4]The complete list of industries is presented in Table 1

Second, we link R&D investment to patent output to measure innovation efficiency—the number of patents generated per unit of R&D expenditure. In this analysis, we construct a probabilistic crosswalk from IPC 4-digit patent classes to ISIC divisions via NAICS industries. Using the IPC-NAICS co-occurrence matrix and official NAICS-to-ISIC crosswalks, we calculate $P(\text{ISIC}|\text{IPC})$ - the probability that a patent in a given IPC class is associated with each ISIC division. We then classify each ISIC division as high-tech, mid-tech, or low-tech following the taxonomy of Fuest et al. (2024) (Table 1), enabling us to aggregate both R&D spending and fractionally-assigned patent counts to a common technology-complexity dimension. We calculate relative innovation efficiency by normalizing each country's patents-per-R&D ratio by the global average for the corresponding technology level, enabling us to identify which sectors and countries translate R&D investment into patentable innovation most effectively. To account for the lag between R&D investment and patent grants, we analyze R&D data from 2015-2019 against patent data from 2018–2020, reflecting an approximate three-year innovation-to-patent cycle.

In both analyses, we use Revealed Comparative Advantage (RCA) to identify technology domains where countries exhibit specialization relative to the global distribution. For country c and IPC class i, we calculate $\text{RCA}_{ci} = (P_{ci}/P_c)/(P_i/P)$, where P_{ci} denotes country c's patent family count in class i, P_c is the country's total patents, P_i is global patents in class i, and P is total global patents. This approach allows us to identify the number of technology fields in which Japan holds comparative advantage within each complexity tier, providing a measure of technological specialization patterns.

Stylized Facts

Despite its reputation as a technological powerhouse, Japan's labor productivity has stagnated over the past two decades, widening the gap with other advanced economies. This paradox is particularly striking given Japan's high levels of R&D investment and industrial sophistication. As Bahar et al. (2024) emphasize, Japan remains one of the most complex economies globally, yet this complexity has not translated into sustained productivity gains.

Figure 1 summarizes the average annual growth rate of labor productivity between 2000 and 2021 across major advanced economies. Japan stands out as the only country with negative productivity growth (although very small in magnitude) during this period, recording an average contraction of 0.19 percent per year. In stark contrast, the United States achieved average gains of 1.22 percent annually, while Germany, France, and the United Kingdom all experienced steady, though moderate, growth ranging from 0.5 to

0.8 percent per year.

This comparison highlights the depth and persistence of Japan's productivity malaise. Over two decades, the cumulative effect of a small negative growth rate implies a significant erosion of relative efficiency—particularly when juxtaposed against peers that continued to compound annual gains. The data reinforce that Japan's stagnation is not merely slower growth but absolute decline, signaling deeper structural inefficiencies.

Figure 1: Average Annual Growth Rate of Labor Productivity for Selected Economies

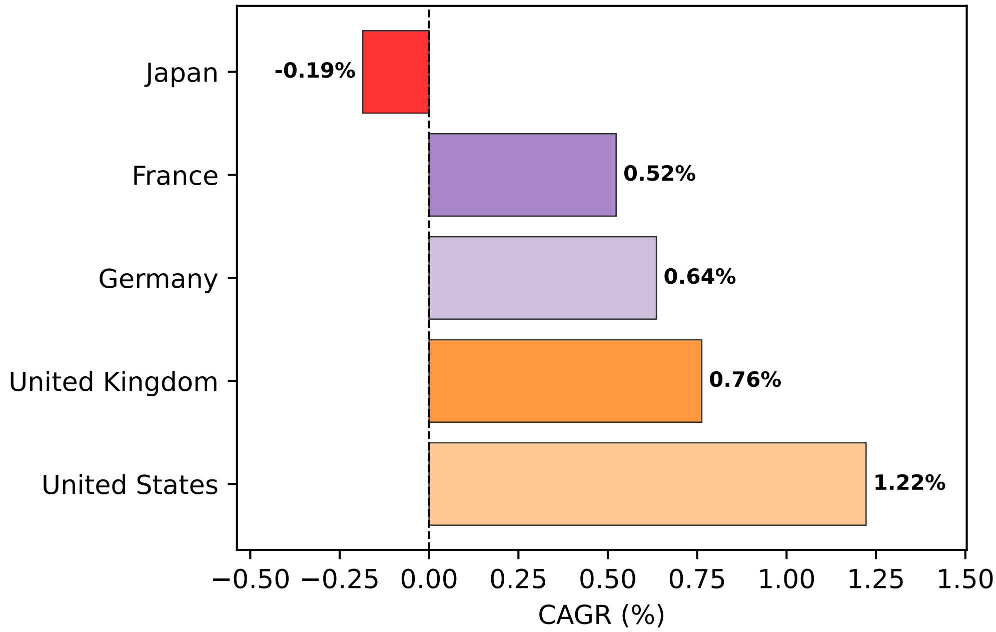

Note: The figure shows the compound annual growth rate (CAGR) of labor productivity for some selected advanced economies between 2000 and 2021.

This important difference in average growth rates evidently results in divergence between the countries in this set of economies. This is shown in Figure 2 which illustrates the evolution of the labor productivity (indexed at 100 in year 2000) for Japan and selected advanced economies between 2000 and 2021. The data reveal a divergence: while countries such as the United States, Germany, France, and the United Kingdom all experienced significant productivity gains—ranging from roughly 15 to 30 percent over the period—Japan's productivity has remained largely flat or even declined. By 2021, Japan's index stood only marginally above its level two decades earlier, underscoring the long-term stagnation in output per worker. The widening productivity gap has major implications for Japan's competitiveness and living standards. This divergence reflects not only slower capital deepening but also weaker total factor productivity growth (TFP) (Bahar and Strauss, 2020).[5]

[5]Recent work has emphasized that some widely discussed cross-country growth gaps shrink consider-

Figure 2: Evolution of Labor Productivity for Japan and Selected Economies

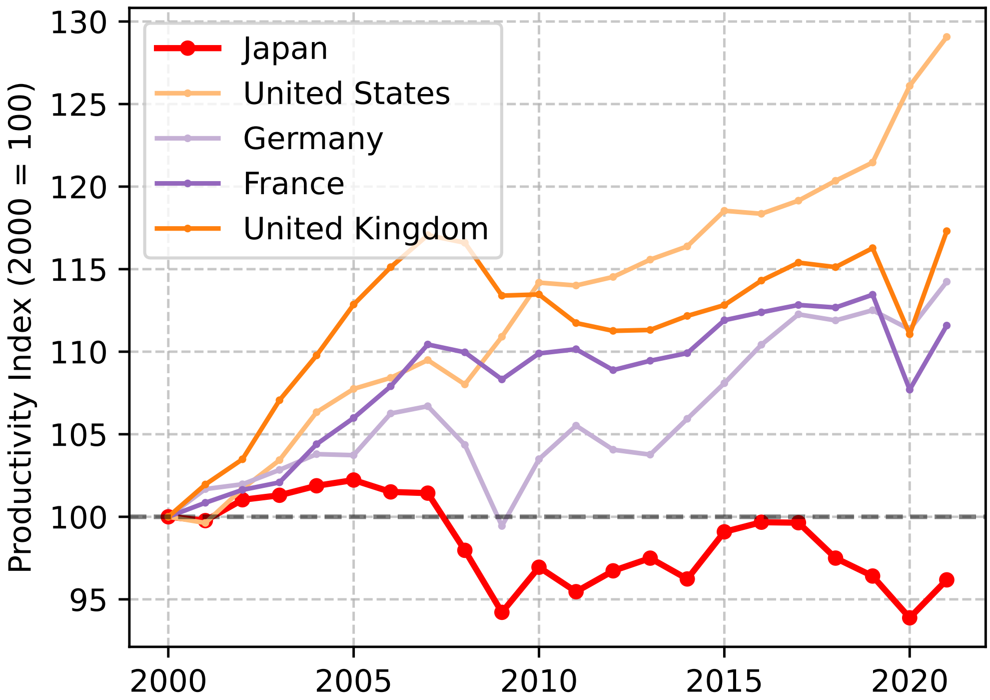

Note: The figure shows the evolution (*i.e.*, growth rate) of labor productivity for some selected advanced economies between 2000 and 2021. The values are indexed to year 2000.

This stagnation is not the result of cyclical downturns alone. Although Japan's productivity dipped sharply following the 2008–2009 global financial crisis, it never recovered to the pre-crisis growth trajectory. In contrast, productivity in peer economies rebounded relatively quickly, maintaining a steady upward trend through the 2010s. The persistent flatlining of Japan's productivity suggests structural rather than transitory factors—consistent with the hypothesis that the country's innovation system has failed to translate technological advances into aggregate efficiency gains.

Figure 3 decomposes Japan's labor productivity dynamics into its main components—gross value added (GVA) and employment—indexed to the year 2000. The chart reveals that Japan's weak productivity performance over the past two decades stems primarily from sluggish growth in value creation rather than from labor-market dynamics. While employment levels (red line) increased steadily after 2012, GVA (blue line) didn't improve, failing to generate proportional gains in output. As a result, the productivity index (green line), defined as GVA per worker, has trended downward since the mid-2000s.

---

ably once one compares GDP per *worker* rather than GDP per capita. For example, Fernández-Villaverde et al. (2023) show that adjusting for demographic structure narrows many apparent differences in economic performance across advanced economies. However, this demographic correction does not materially alter the picture for Japan, we claim. Even when focusing on GDP or GVA per worker—rather than per capita—Japan's productivity performance remains stagnant relative to its peers, and its TFP growth continues to lag behind other advanced economies. In this sense, the stagnation highlighted in Figure 2 reflects genuine productivity weakness rather than a purely demographic artifact.

In most advanced economies, productivity growth typically arises from expansions in GVA that outpace employment growth, reflecting technological progress, capital deepening, and efficiency improvements. In Japan, however, the opposite pattern emerges: employment has grown faster than output, implying diminishing returns to labor and limited gains in efficiency. This suggests that Japan's firms are not creating new value at the pace required to sustain productivity growth. Instead, output growth has been largely absorbed by rising labor input, leaving per-worker productivity stagnant or declining.

Figure 3: Japan's Evolution of Labor Productivity and its Components

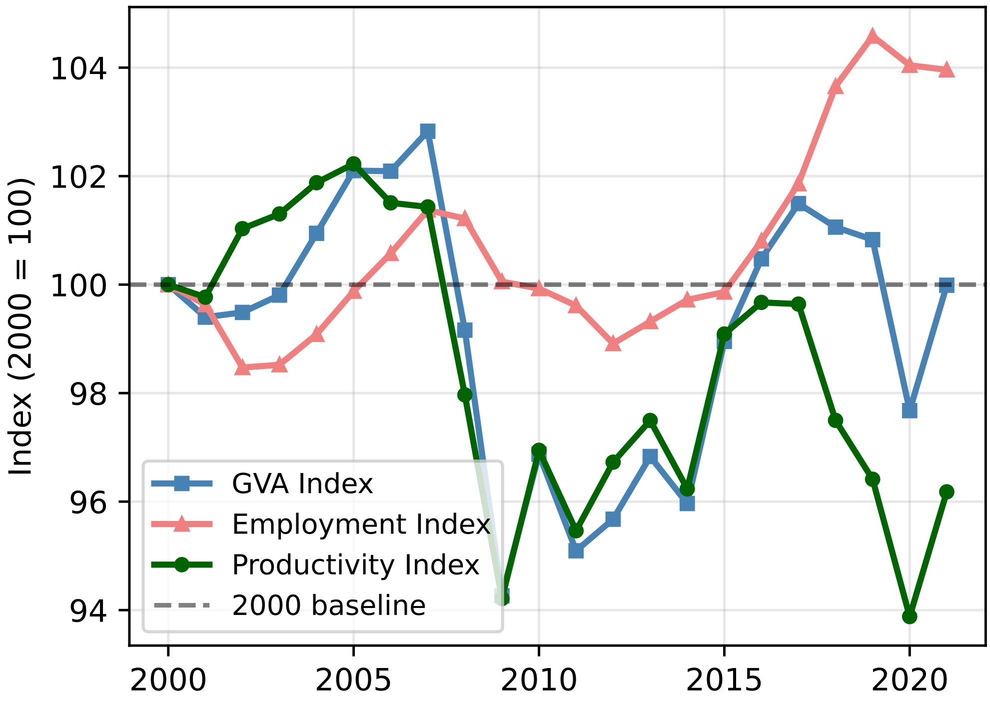

Note: The figure shows the evolution (*i.e.*, growth rate) of labor productivity and its components for Japan between 2000 and 2021. The values are indexed to year 2000.

As such, the creation of value added remains a central issue in Japan's economic performance. As we demonstrated in our earlier paper (Bahar et al., 2024), Japanese firms have increasingly chosen to expand their most profitable activities abroad. This strategic internationalization—through foreign direct investment and intellectual property licensing—has allowed them to sustain high returns in the face of mounting domestic challenges, including demographic decline, labor scarcity, and structural rigidities. Yet, while this global redeployment of Japan's productive knowledge has been highly successful, it has also left the domestic economy struggling to generate new sources of value. The natural question, then, is about Japan's capacity to create value added at home, focusing specifically on how innovation dynamics—its allocation, composition, and quality—shape the country's productivity outcomes.

Focusing on Innovation

Figure 4 compares R&D intensity—measured as research and development expenditure as a share of GVA—across the top fifteen economies in 2009 and 2021.[6] Japan stands among the world's leaders, maintaining a consistently high R&D intensity of around 2.6 percent of GVA in 2021, up slightly from 2.4 percent in 2009. This places Japan well above major European economies such as Germany, and nearly on par with the United States. Only South Korea significantly outpaces Japan in our sample, with R&D spending exceeding 4 percent of GVA by 2021.

While these figures underscore Japan's continued commitment to technological advancement, they also highlight a paradox. Despite its high and sustained investment in R&D, Japan's aggregate productivity growth has remained stagnant, as shown in the preceding section. This suggests that Japan's challenge is not one of innovation *effort* but of innovation effectiveness—that is, the ability to convert R&D inputs into tangible productivity gains and value creation.

Figure 4: R&D Intensity for Selected Economies

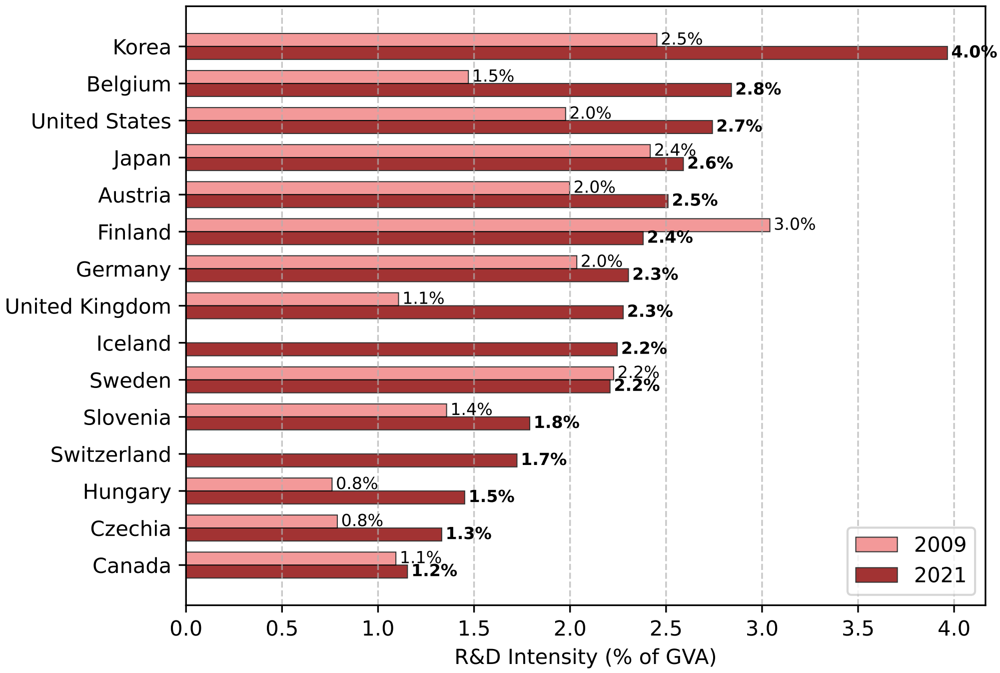

Note: The figure shows R&D Intensity–R&D expenditures as share of Gross Value Added–for the top spending economies in our dataset for 2009 and 2021.

To unpack this, it is essential to look beyond aggregate R&D intensity and examine how R&D resources are allocated across industries—and how effective those investments

[6]As mentioned in the data section of this document, we use a panel of countries with industry-level data from 2000 to 2021. However, in this panel we only have R&D expenditures data from 2009 onward. Because of this, the innovation analysis that involves R&D data has a different time window compared to the analysis on productivity.

are in generating value added. In other words, understanding Japan's productivity stagnation requires not only measuring *how much* the country invests in innovation, but also *where* those investments go and *what they yield*. By studying the distribution of R&D across sectors with different technological intensities, we can begin to assess whether Japan's innovation system is channeling resources toward activities that expand the technological frontier or instead reinforcing incremental improvements in mature industries.

Figures 5 and 6 illustrate the structural composition of Japan's private R&D spending compared to other advanced economies. Panel (a) shows that Japan allocates an exceptionally high share—nearly 87 percent—of its private R&D investment to the secondary sector, which includes manufacturing, utilities, and construction. This stands in sharp contrast to the United States, where only 55 percent of R&D is concentrated in the secondary sector and nearly 45 percent occurs in services. Other high-income economies, such as Germany (83 percent) and Switzerland (91 percent), also maintain strong manufacturing R&D bases, yet Japan's profile remains particularly skewed toward manufacturing and away from services, where much of the global technological frontier has shifted in recent decades.

Figure 5: R&D Composition by Economic Sector

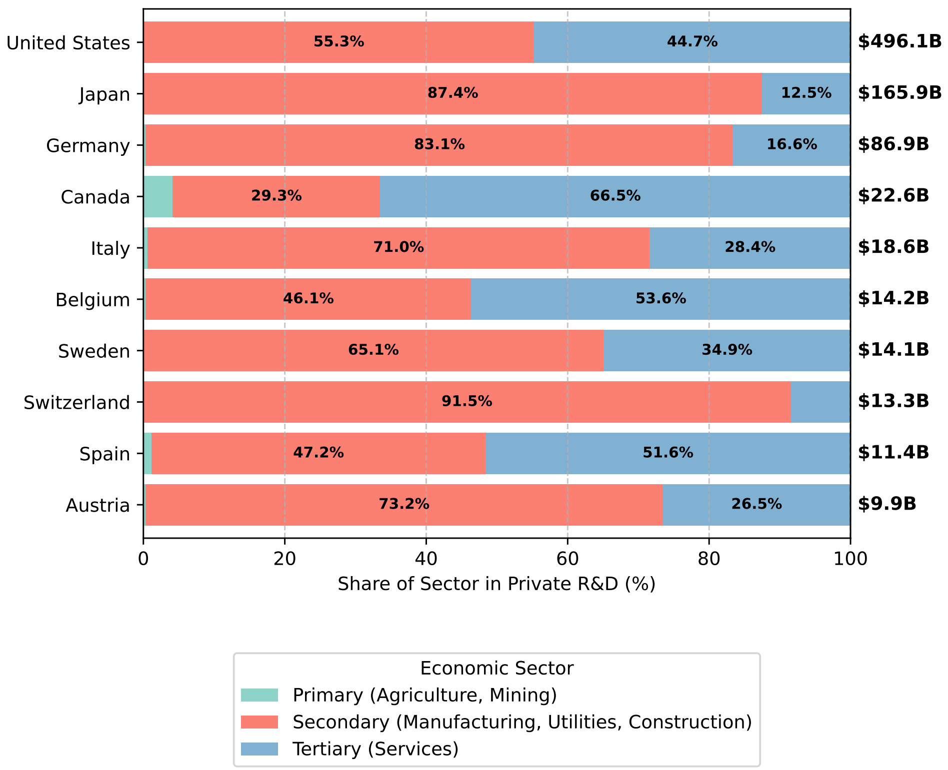

Note: The figure shows the distribution of total private R&D expenditures across the primary, secondary, and tertiary sectors for the top R&D-spending countries in 2021. Bars indicate each sector's share of national R&D, while the bold values on the right report total R&D in real US dollars (2010 base year).

Figure 6: R&D Composition in Manufacturing

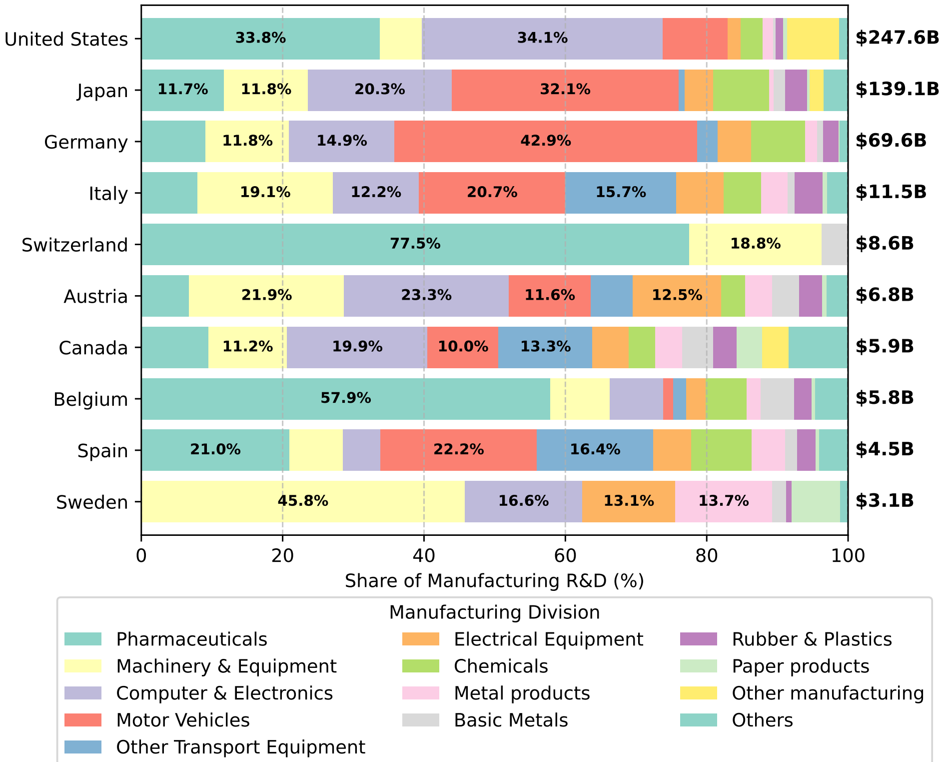

Note: The figure displays the composition of Manufacturing R&D for top-spending countries in 2021. Bars show the share of total manufacturing R&D attributed to each division, and the bold values on the right give the corresponding R&D amounts in real US dollars (2010 base year).

This composition underscores one of the defining features of Japan's innovation system: an enduring emphasis on industrial and engineering excellence, especially in traditional manufacturing domains. While this focus has historically fueled Japan's global competitiveness, it also raises questions about its adaptability to an increasingly service- and knowledge-intensive economy. For instance, in 2021 only about 4.5% of Japan's total R&D expenditures were allocated to the tertiary sector, and within that, nearly 80% was concentrated in scientific research and development activities. By contrast, in the United States, a much larger share of tertiary R&D—almost 45%—was directed toward software-related and digital activities (Figure 7). In these sectors, where R&D in advanced economies has grown rapidly, Japan remains relatively underinvested. This structural bias may explain why Japan's aggregate R&D intensity is high but its productivity growth limited: most of its innovation occurs in sectors that face diminishing returns to incremental technological improvement.

Figure 7: Composition of R&D Expenditures in Services

Note: The figure displays the distribution of R&D expenditures within the tertiary (services) sector across selected top-spending economies in 2021. Bars show the share of total service-sector R&D attributed to each division, while the bold values on the right indicate the corresponding R&D amounts in real US dollars (2010 base year).

In this context, we build on the framework recently proposed by Fuest et al. (2024), who introduced the *"middle-technology trap"* hypothesis in their analysis of EU innovation policy. Their central argument is that economies heavily concentrated in middle-technology industries—technologically sophisticated but primarily applying advanced technologies developed elsewhere rather than originating new ones—may fall into a structurally low-growth equilibrium, in which they contribute less to overall productivity growth despite absorbing a large share of R&D investment. Our contribution is to test this hypothesis empirically and at scale, extending the analysis from their EU/US focus to Japan, benchmarking Japan against a broad set of advanced economies using tangible measures of innovation—R&D inputs and patent outputs—and demonstrating robustness across nine alternative industry classifications. Applying this framework to Japan allows us to test whether its pattern of R&D concentration—particularly in industries such as automotive and electronics manufacturing—has constrained its ability to generate new value and move the productivity frontier forward.

We build on the framework of Fuest et al. (2024) to develop an industry-level technology classification and examine how the structure of R&D investment shapes productivity outcomes. Their taxonomy assigns each industry to one of three broad categories—high-tech, mid-tech, or other.[7] The core principle for distinguishing high-tech and mid-tech industries reflects the idea that mid-tech industries might use advanced technologies as inputs (e.g., autonomous driving software in vehicles), but these technologies are created in "upstream" sectors (high-tech). All remaining sectors are classified by the authors as "other". Using this taxonomy, we assign each ISIC industry in our sample to a technology category, where the industries defined as "other" by Fuest et al. (2024) correspond to low-tech in our classification.

As illustrated in Table 1, high-tech sectors include industries such as computer, electronic and optical products (C26) and information and communication (J), where innovation typically involves frontier technologies and generates strong spillovers across the economy. In contrast, mid-tech industries—including chemicals (C20), electrical equipment (C27), and especially motor vehicles and transport equipment (C29–C30)—rely on advanced technological inputs but tend to focus on incremental innovation and process improvements rather than on the creation of entirely new technologies. The remaining industries encompass lower-tech manufacturing, services, and non-tradable activities with

[7] Specifically, Fuest et al. (2024) use the following classification:

- **High-tech**: aerospace & defence, alternative energy, electronic & electrical equipment, health care equipment & services, pharmaceuticals & biotechnology, software & computer services, and technology hardware & equipment.
- **Mid-tech**: automobiles & parts, chemicals, financial services, fixed line telecommunications, industrial engineering, industrial metals & mining, industrial transportation, leisure goods, mobile telecommunications, and personal goods.
- **Other**: banks, beverages, construction & materials, electricity, food & drug retailers, food producers, forestry & paper, gas, water & multiutilities, general industrials, general retailers, household goods & home construction, life insurance, media, mining, nonlife insurance, oil & gas producers, oil equipment, services & distribution, real estate investment & services, support services, tobacco, and travel & leisure.

limited innovation intensity.

Table 1: Industry Classification by Technology Level based on Fuest *et al.* Framework

ISIC Code	ISIC Level	Description	Low-tech	Mid-tech	High-tech
A01	division	Crop and animal production, hunting and related service activities	X		
A02	division	Forestry and logging	X		
B	section	Mining and quarrying		X	
C16	division	Manufacture of wood and of products of wood and cork	X		
C17	division	Manufacture of paper and paper products	X		
C18	division	Printing and reproduction of recorded media	X		
C19	division	Manufacture of coke and refined petroleum products		X	
C20	division	Manufacture of chemicals and chemical products		X	
C21	division	Manufacture of basic pharmaceuticals			X
C22	division	Manufacture of rubber and plastics products		X	
C23	division	Manufacture of other non-metallic mineral products		X	
C24	division	Manufacture of basic metals		X	
C25	division	Manufacture of fabricated metal products, except machinery	X		
C26	division	Manufacture of computer, electronic and optical products			X
C27	division	Manufacture of electrical equipment			X
C28	division	Manufacture of machinery and equipment n.e.c.		X	
C29	division	Manufacture of motor vehicles, trailers and semi-trailers		X	
C30	division	Manufacture of other transport equipment			X
F	section	Construction	X		
G46	division	Wholesale trade, except of motor vehicles and motorcycles	X		
G47	division	Retail trade, except of motor vehicles and motorcycles	X		
H49	division	Land transport and transport via pipelines		X	
H50	division	Water transport		X	
H51	division	Air transport		X	
H52	division	Warehousing and support activities for transportation		X	
H53	division	Postal and courier activities	X		
I	section	Accommodation and food service activities	X		
J	section	Information and communication			X
K64	division	Financial service activities, except insurance and pension funding	X		
K65	division	Insurance, reinsurance and pension funding	X		
L	section	Real estate activities	X		
M	section	Professional, scientific and technical activities	X		
N77	division	Rental and leasing activities	X		
O84	division	Public Administration and Defence	X		
P85	division	Education	X		
Q86	division	Human health activities	X		
R	section	Arts, entertainment and recreation	X		

Using this classification, Figures 8a and 8b provide a comparative overview of the composition of R&D expenditure by technology level across regions. Figure 8a shows the distribution of total R&D spending in 2021 by technology class. The pattern is striking: Japan allocates more than half (53.9 percent) of its total R&D to mid-technology industries, a markedly higher share than the United States (18.4 percent) or the European Union (42.4 percent). In contrast, high-technology industries account for only about 39 percent of Japan's R&D, compared to 69 percent in the United States and 50 percent in the EU. This composition highlights Japan's structural bias toward technologically advanced but mature manufacturing sectors—such as motor vehicles, electrical equipment, and chemicals—rather than frontier, high-tech domains like semiconductors, biotechnology, or digital services.

Figure 8b traces the evolution of mid-tech R&D shares between 2009 and 2021. Japan's share of mid-technology R&D has risen steadily, surpassing 50 percent in recent years, even as the corresponding share declined or stabilized in other major economies. The United States and China both exhibit relatively low and decreasing shares of mid-

Figure 8: Composition of R&D Expenditure by Technology Level Across Regions

(a): R&D by Technology Level, 2021 (b): R&D Share in Mid-Tech (2009–2021)

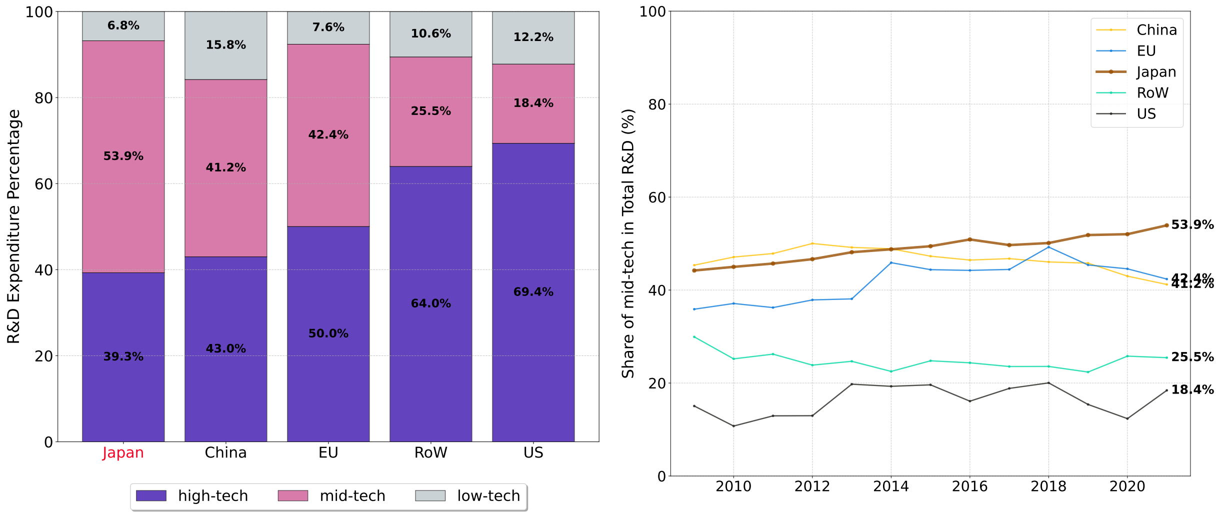

Note: The figure shows the composition of R&D expenditures by technology level across regions. Panel (a) reports the 2021 distribution, while panel (b) displays the evolution of the mid-tech R&D share from 2009 to 2021.

tech R&D, while the European Union and the rest of the world maintain moderate levels. Japan thus stands out not only for the scale but also for the persistence of its commitment to mid-tech innovation.

While Figure 8 demonstrates that Japan channels the majority of its R&D expenditures toward mid-technology industries, it is important to consider the sensitivity of this finding to the specific classification scheme used. One potential limitation of this analysis is that our results may be driven by how industries are categorized into high-tech, mid-tech, or low-tech sectors. To address this concern and test the robustness of our findings, we developed and applied eight alternative classification methods in addition to the taxonomy of Fuest et al. (2024). A summary of these classifications is presented below.[8]

1. **Average patents (avg_patents):** Measures the average volume of patent activity per technology class (2015–2021). Technology classes with higher patent volumes are assigned to higher-tech categories, reflecting greater innovation intensity.

2. **Innovation Momentum (IM):** Captures the compound annual growth rate (CAGR) of global patent families within each technology class from 2015 to 2021. Positive

[8]We note that the metrics and terminology presented here—including Patent Flow Recency, Innovation Momentum, and Structural Age—are terms used for convenience in this analysis and are not standard terms from the existing literature. We introduce them for convenience and to capture distinct dimensions of technological dynamism.

values indicate expanding technological domains, while negative values suggest declining or mature fields. Higher IM scores indicate technologies with sustained growth momentum.

3. **Median Age (MA):** Measures the median age of patents within each technology class, calculated as the time elapsed since patent filing. Lower median ages indicate more recent innovation activity and emerging technologies, while higher values suggest mature or declining fields.

4. **Patent Flow Recency (PFR):** Measures the share of patents within a technology class that were filed in recent years (2019–2021). Values range from 0 to 1, with higher scores indicating concentration of recent innovation activity, either reflecting genuine technological emergence or significant renewal within an existing domain.

5. **Product Complexity Index (naics_pci):** We use Product Complexity Index values calculated from US County Business Patterns data using the methodology of Hausmann et al. (2014). These PCI values measure the complexity of economic activities at the NAICS 6-digit industry level, calculated from employment counts across US counties.

6. **LLM-based metrics:** These classifications use large language model (LLM) assessments to assign industries to technological categories.

   - The *llm_inputs* metric classifies industries according to the technological sophistication of their inputs, such as required workforce skills or dependence on advanced processes.
   - The *llm_outputs* metric classifies industries based on the technological advancement embodied in their outputs, considering the novelty, complexity, and innovation reflected in products, services, or knowledge produced.
   - The *llm_holistic* metric considers both input-side characteristics (*e.g.*, skills, technology, and processes) and output-side features (*e.g.*, product complexity and innovation) to assess each industry's overall technological level.

As shown in Figure 9, Japan's concentration of R&D in mid-tech industries remains highly robust across these various classification schemes. Across nearly all these frameworks, mid-tech sectors consistently account for between 50 and 70 percent of Japan's total R&D, while high-tech industries rarely exceed 40 percent.[9]

[9]Based on the values produced by the alternative classification metrics, we constructed a consensus classification using a majority voting mechanism. For the remainder of this paper, we use this consensus-based technology classification unless otherwise noted.

Figure 9: Robustness of Japan's mid-tech R&D Concentration Across Alternative Classification Schemes

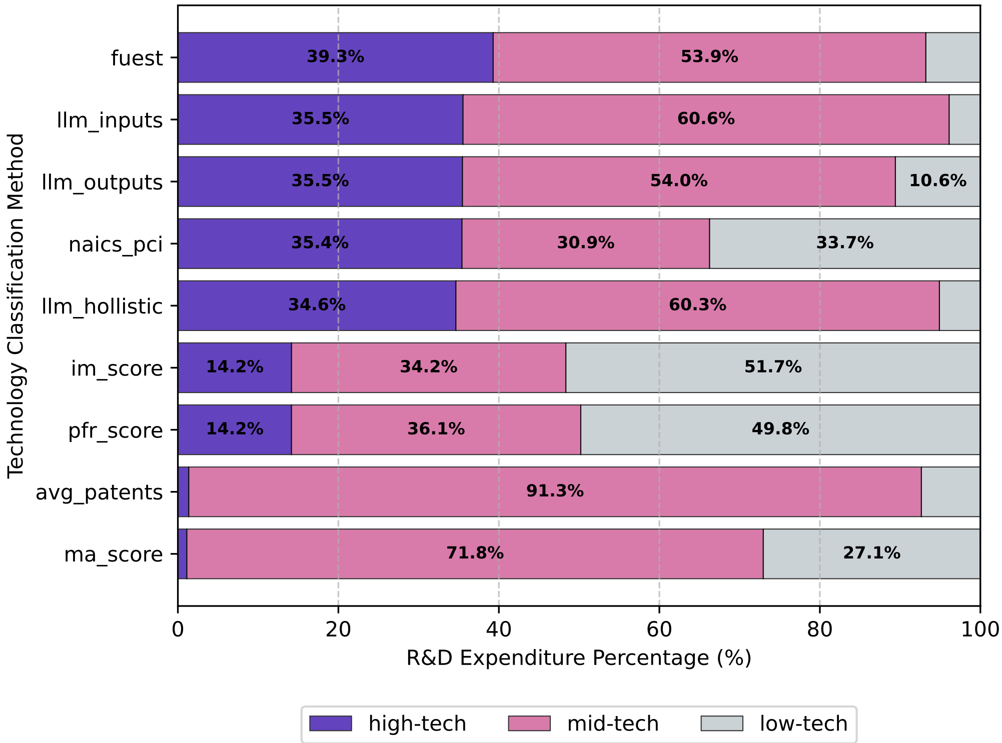

Note: This figure compares Japan's 2021 R&D allocation across high-, mid-, and low-technology industries using nine alternative classification schemes, illustrating that the distribution is stable across diverse ways of assigning industries into technological classifications.

This pattern of R&D concentration is reflected, though imperfectly, in Japan's patent output. Figure 10 presents the composition of patent output by technology level for Japan and selected comparison countries for the period 2018–2020. Japan's patent portfolio shows 46.4 percent in high-technology domains, 39.7 percent in mid-technology, and 13.9 percent in low-technology sectors. While Japan allocates more than half of its R&D expenditure to mid-technology industries (as shown in Figure 8a), only 40 percent of its patent output falls into this category. This gap between R&D investment concentration (54 percent) and patent output share (40 percent) in mid-tech sectors suggests that Japan's innovation system invests heavily in these domains but generates intellectual property at a lower rate than its R&D allocation would predict.

Figure 10: Patent Output Composition by Technology Level Across Major Economies (2018–2020)

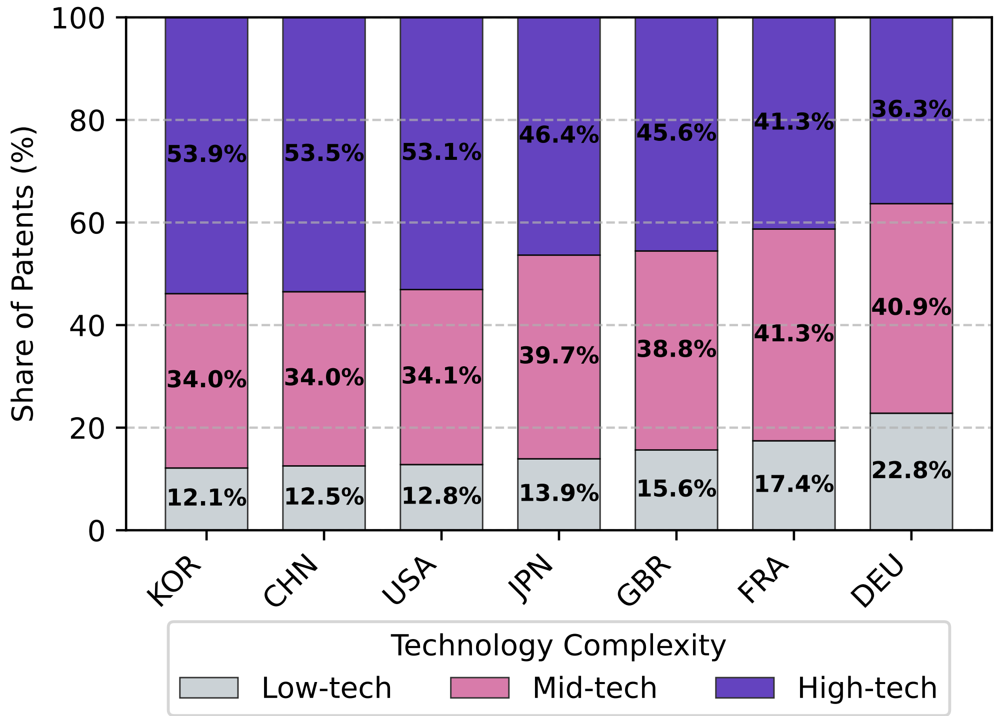

Comparing Japan to peer economies reveals that major innovation leaders-including the United States, South Korea, and even China-exhibit substantially higher shares of patents in high-technology sectors, typically exceeding 44 percent. This pattern persists when we examine comparative advantage rather than raw patent volumes (Figure 11). We calculate revealed comparative advantage (RCA > 1) across technology domains and find that only 19.3 percent of Japan's comparative advantages lie in high-tech sectors, compared to 55.1 percent in mid-tech and 25.6 percent in low-tech. This distribution reinforces that Japan's relative strengths are tilted away from high-technology domains and toward mid- and low-technology sectors.

Taken together, these figures reinforce the conclusion that Japan's innovation activ-

Figure 11: Distribution of Comparative Advantages (RCA > 1) by Technology Level (2018–2020)

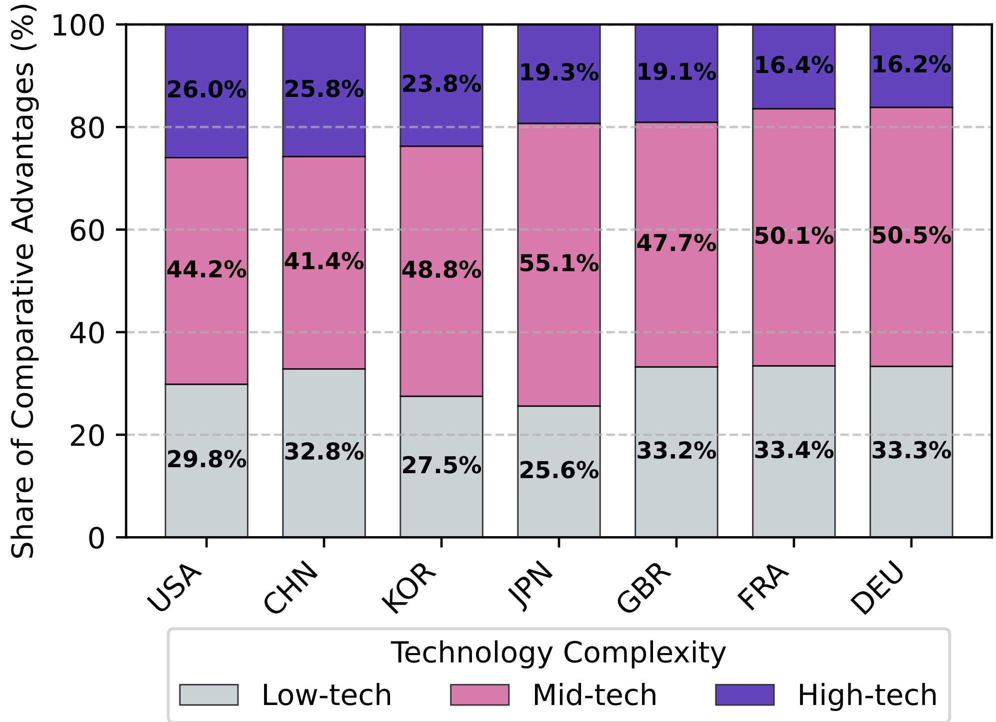

ity remains heavily concentrated in mid-technology industries. While this concentration may support competitiveness in established sectors, it may also constrain the economy's capacity to generate new sources of growth. In the following section, we use the taxonomy presented here to decompose Japan's productivity growth and examine how the contributions of high-, mid-, and low-technology industries differ across economies. We also assess whether Japan's innovation orientation toward mid-tech industries helps to explain the country's persistent gap between innovation effort and productivity outcomes.

A Bug or a Feature?

While our analysis so far has documented Japan's disproportionate concentration of R&D in mid-technology sectors, it is important to ask whether this pattern represents a bug or a feature of Japan's innovation model. In other words, is the strong emphasis on mid-tech industries—such as automotive, electrical equipment, and chemicals—a rational specialization strategy that effectively supports growth under Japan's constraints, or is it a structural distortion that limits the country's capacity to generate productivity gains?

Figure 12 plots the compound annual growth rate (CAGR)[10] of labor productivity

[10]We use CAGR as it is the geometric version of the average annual growth rate, which is a more accurate reflection of consistent long-term growth.

by technology level for a group of high-income countries between 2000 and 2021. It shows that in the United States high-technology industries have recorded the largest productivity growth, reflecting their role in driving technological progress and value creation. In Germany, the story is less clear, though high-tech industries productivity growth has grown somewhat. In Japan, however, the pattern is reversed: productivity growth in high- and mid-tech industries has been negative. This evidence suggests that Japan's innovation model, while intensive, is not translating into efficiency improvements where it matters most. The results underscore that Japan's focus on mid-tech industries may be a feature of its industrial legacy but a bug for its long-term productivity performance, as innovation efforts in these sectors have yielded limited aggregate returns.

Figure 12: Average Annual Productivity Growth Rate (CAGR) by Technology Level for Selected Countries

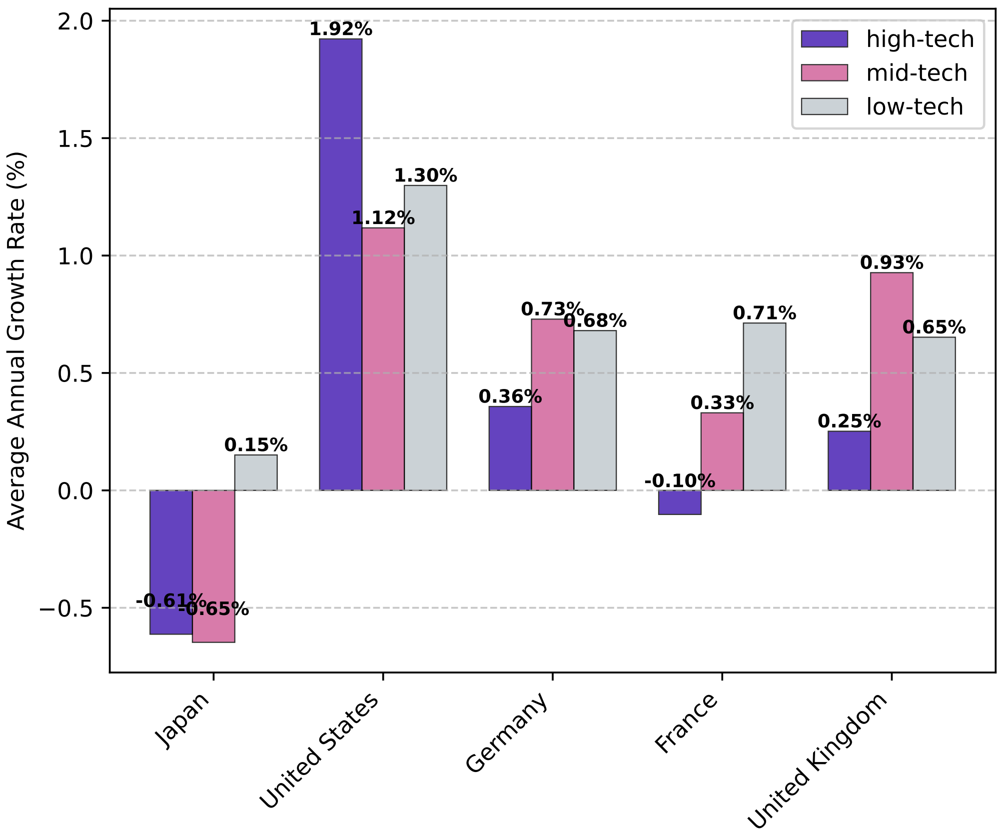

Note: This figure compares average annual productivity growth from 2000 to 2021 across high-, mid-, and low-technology industries for selected economies, illustrating how productivity performance varies both across countries and across technology tiers.

To assess whether different types of industries—high-tech, mid-tech, or low tech—are systematically associated with higher productivity growth, we conduct a regression analysis using a cross-country, industry-level dataset. The goal is to determine, on average, which types of industries contribute the most to aggregate productivity growth, and whether Japan's emphasis on mid-technology sectors aligns with this broader global pattern.

Our empirical specification takes the form:

$$Y_{i,j} = \beta_0 + \beta_1 \text{HighTech}_i + \beta_2 \text{MidTech}_i + \beta_3 \text{GVAShare}_{i,j} + \gamma_j + \varepsilon_{i,j} \tag{1}$$

The variable $Y_{i,j}$ denotes the contribution of sector i to the annual average growth of labor productivity in nation j. The variables HighTech_{ij} and MidTech_{ij} function as binary indicators to determine whether the industry is categorized within each technological group, as per the classifications previously delineated. The variable $\text{GVAShare}_{i,j}$ serves as a control variable, representing the proportionate significance of industry i within the economy. Furthermore, the parameter γ_j accounts for country-specific fixed effects, thus controlling for macroeconomic and institutional disparities across different countries. The dataset employed in this analysis encompasses the same nations and industries detailed in the data section. Notably, to evaluate long-term effects, we compute the contribution of industry i to productivity growth in country j across the entire period from 2000 to 2021. Consequently, the regression analysis utilizes 885 observations.

Figure 13 summarizes the estimated coefficients from our regression analysis, showing the cumulative technology effects on total productivity contribution across multiple classification schemes. Each bar represents the average effect size (in percentage points) for industries classified as low-tech (baseline in gray), baseline and mid-tech (pink), and baseline and high-tech (dark purple), while the different groups along the x-axis correspond to the alternative classification methods described in the previous section.

The results are remarkably consistent across all schemes. On average, high-technology industries exhibit the strongest positive association with productivity growth, contributing between 0.07 and 0.10 percentage points more than low-technology industries. By contrast, mid-technology industries contribute only marginally—often less than half the high-tech effect—and in some classifications, their estimated impact is statistically indistinguishable from zero.

These results confirm that the link between technological sophistication and productivity growth is robust: economies with large shares of high-tech activity are more likely experience stronger productivity performance. Importantly, this evidence suggests that Japan's concentration of R&D and innovation effort in mid-tech sectors is not an optimal configuration from a productivity standpoint. Rather, it points to a structural inefficiency, in which significant innovation resources are devoted to industries that, on average, yield smaller aggregate returns. This is something we explore next.

We begin by calculating an allocation metric that captures how intensively Japan directs research spending toward a given industry relative to that industry's weight in

Figure 13: Estimated Technology Effects on Productivity Contribution Across Classification Schemes

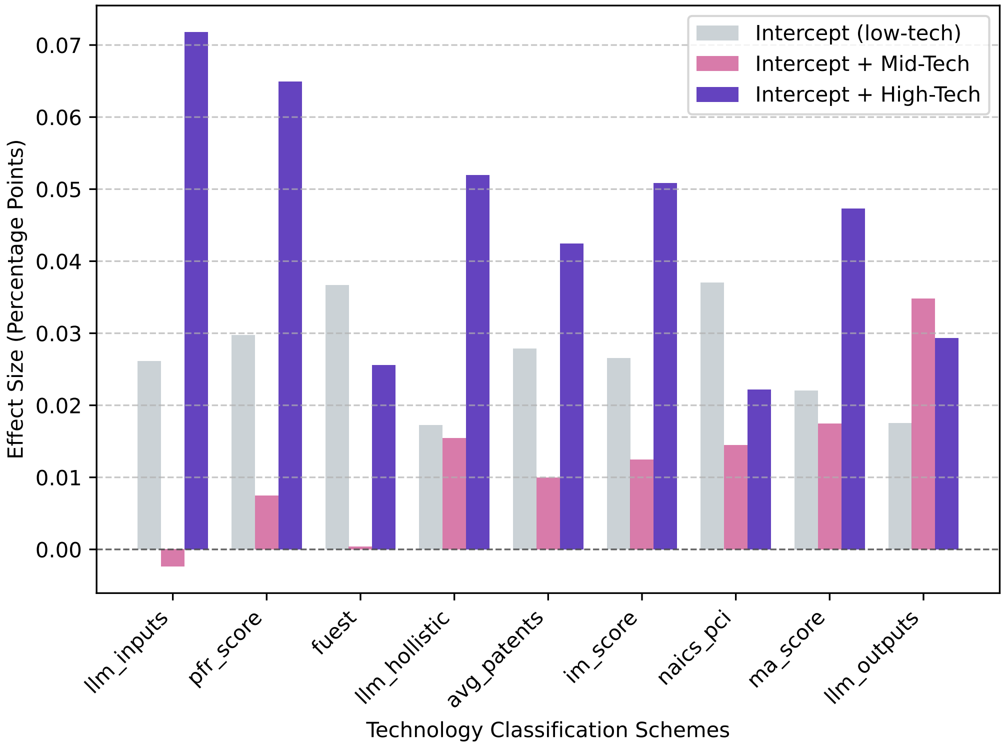

Note: This figure shows the cumulative technology effects on total productivity contribution estimated from the regression analysis across nine alternative classification schemes. For each scheme, the gray bar represents the baseline effect for low-tech industries, while the pink and dark-purple bars add the incremental effects for mid-tech and high-tech industries, respectively. The height of each bar therefore reflects the average percentage-point contribution associated with belonging to that technology tier under the corresponding classification method.

the domestic economy.[11] Figure 14 presents this metric relative to the industry sample average. Values greater than one indicate that, relative to the rest of the world, an industry allocated more resources to R&D than the value it added to the economy. Such patterns might signal a potential over-allocation of resources; nonetheless, some industries naturally require sustained research investment because payoffs materialize only after long development cycles. Conversely, values below one could indicate untapped potential for additional investment; however, not every industry requires or benefits from extensive R&D, especially if processes are highly standardized or if competitive advantage depends more on scale, logistics, and cost management than on technological upgrading. In Japan, the industries where R&D allocation appears highest are predominantly mid-technology industries.

Figure 14: R&D Allocation in Japan Relative to Sample Averages

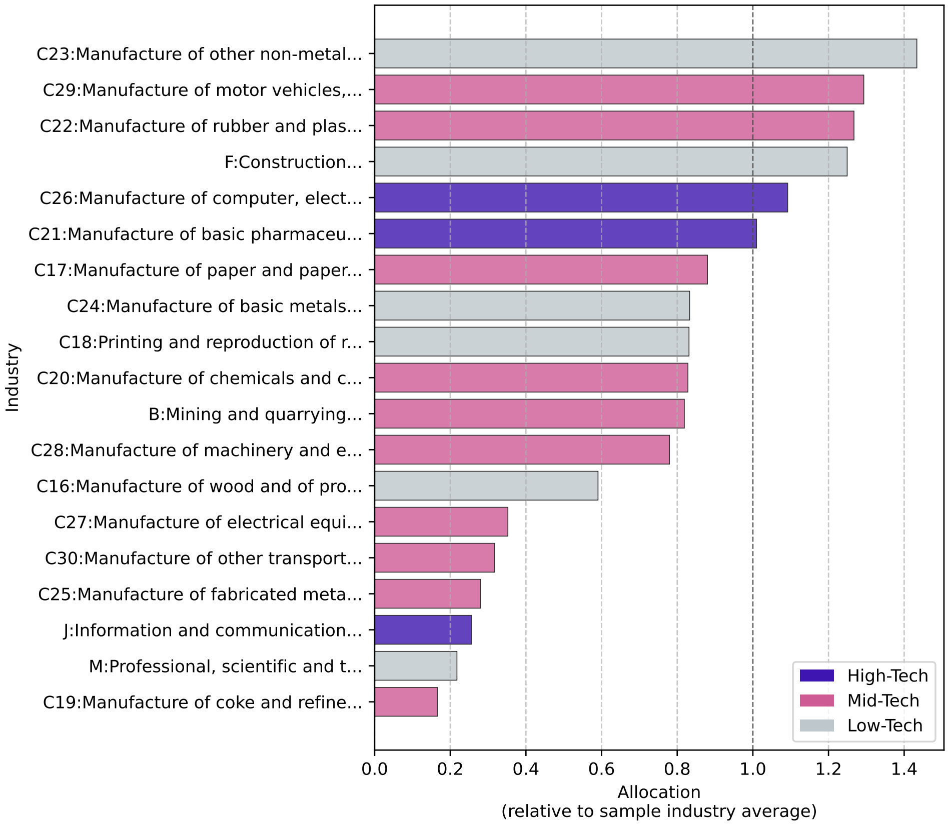

Note: This figure displays Japan's R&D allocation across industries, measured as each industry's ratio of R&D expenditure to its share of domestic value added, relative to the sample-wide industry average. Bars above one indicate industries receiving a disproportionately large share of R&D resources, while values below one indicate comparatively lower R&D intensity.

Having established that mid-technology industries dominate Japan's innovation landscape, it becomes crucial to examine whether this concentration reflects an efficient spe-

[11] Specifically, the allocation metric is the ratio between the average share of R&D expenditure and the average share of GVA for a specific industry between 2009 and 2021.

cialization or a structural inefficiency. In other words, if Japan's focus on mid-tech industries is a feature rather than a bug, one would expect these sectors to exhibit higher returns per unit of R&D investment. Figure 15 directly tests this hypothesis by plotting the previously described allocation metric for each industry (horizontal axis) against its contribution to labor productivity growth (vertical axis), both measured relative to the industry sample average. The size of each bubble corresponds to the industry's share in Japan's GVA.

Figure 15: R&D Allocation and Productivity Contribution for Japanese Industries

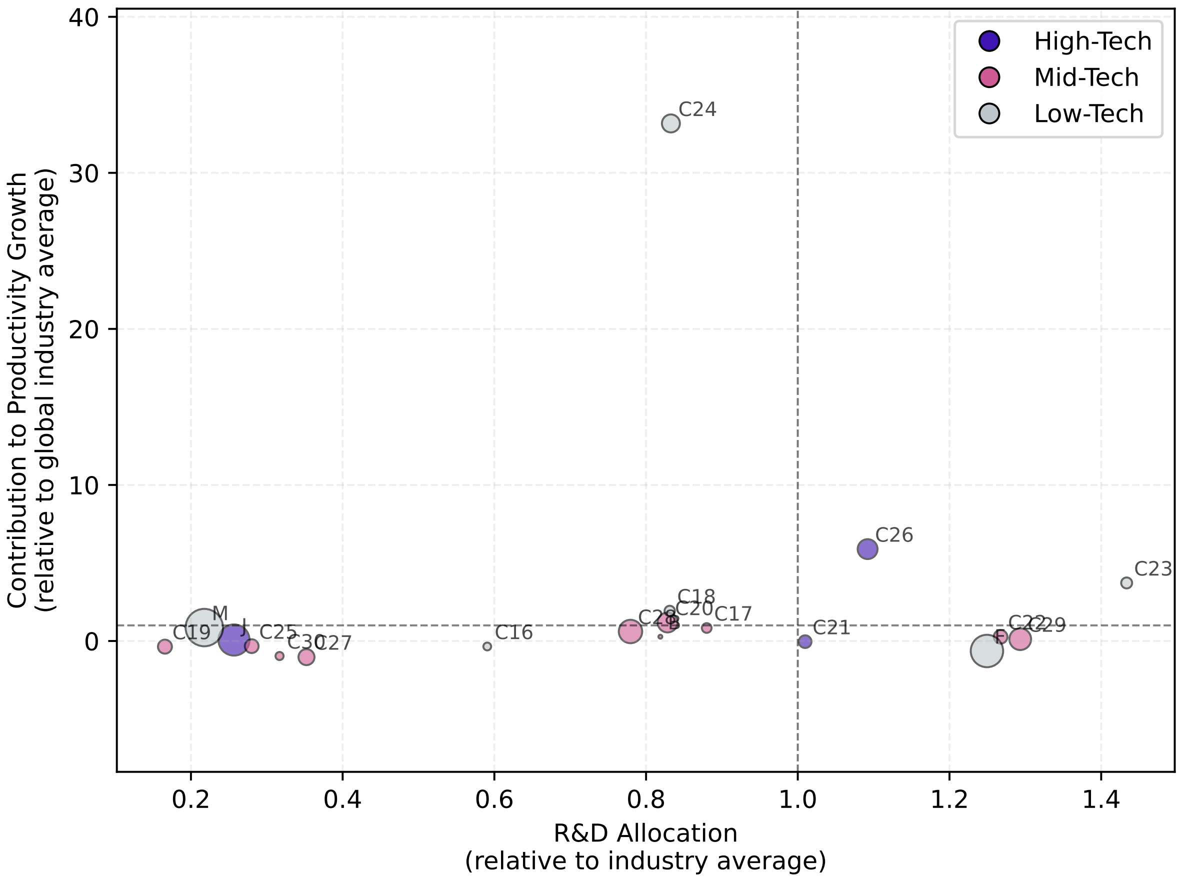

Note: The figure divides industries into four quadrants. The upper-right quadrant represents sectors with above-average R&D allocation and above-average productivity contribution—industries where innovation appears efficient and productive. The lower-right quadrant captures above-average allocation but below-average productivity contribution, indicating potential inefficiencies or over-investment. The left-hand quadrants correspond to industries where Japan invests less than the global average, either appropriately (low allocation and low productivity) or potentially sub-optimally (low allocation despite high productivity payoff).

For Japan, the pattern is striking. Mid-tech industries, such as manufacturing of motor vehicles (C29) and rubber plastics (C22), lie in the lower-right quadrant—sectors where Japan allocates a disproportionate share of R&D resources yet obtains below-average productivity gains compared to their peers. This suggests that Japan's innovation investments in these mature manufacturing industries generate diminishing productivity returns, consistent with the idea of a "middle-technology trap." In contrast, only one high-technology industry (computer, electronic, and optical products, C26) appears in the upper-right quadrant, indicating relatively higher efficiency in translating R&D inputs into productivity growth.

There is, however, one outlier that merits explicit discussion: basic metals (C24). Unlike most mid-technology sectors, this one sits in the upper-left quadrant, meaning it delivers above-average productivity contribution despite below-average R&D allocation. This asymmetric position suggests a sector that generates value-added growth not through intensive innovation spending, but through accumulated technological capabilities, process optimization, and long-standing industrial competitiveness. Japan's steel sector is a natural example. Firms such as Nippon Steel—the world's fourth-largest producer and historically one of the most technologically advanced steelmakers—have achieved productivity gains that exceed what their R&D intensity alone would predict.

The existence of C24 as a high-productivity, low-R&D outlier therefore reinforces the central argument of this section: Japan's innovation challenge is not simply one of spending more, but of spending more where returns are structurally higher. Basic metals demonstrates that some industries continue to generate strong productivity growth even without large, persistent R&D intensity. Most mid-tech manufacturing sectors do not. The contrast between C24 and C29 is illustrative: both are technologically sophisticated and historically strong, but only one still produces measurable productivity dividends.

Overall, the evidence implies that Japan's R&D system may not be optimizing for productivity impact. Rather than delivering superior productivity outcomes in the industries where it invests most heavily, Japan's mid-tech sectors under-perform relative to their global benchmarks. This pattern strengthens the interpretation that the country's innovation specialization—rooted in its industrial legacy and reinforced by institutional and policy incentives—is not yielding the expected returns in terms of value creation. Moving forward, reorienting part of Japan's innovation effort toward high-technology sectors and emerging service industries may be necessary to enhance the efficiency and aggregate impact of its R&D investments.

One potential explanation for Japan's mid-tech concentration could be that these sectors offer higher returns to innovation effort, justifying the allocation pattern. However, Figure 16 shows that this is not the case. The figure compares innovation efficiency - measured as patents generated per million U.S. dollars of R&D investment - across countries and technology levels, relative to the global average for each technology category. Japan's efficiency in converting R&D investment into patent output varies substantially across technology levels. In high-technology sectors, Japan's relative efficiency stands at approximately 1.47, indicating that Japan generates 47 percent more patents per dollar of R&D than the global average. In mid-technology industries, however, Japan's relative efficiency falls to approximately 0.73, meaning Japan produces only 73 percent as many patents per dollar of R&D investment as the global average for mid-tech sectors. Low-tech efficiency is lower still, at approximately 0.56. This pattern suggests that Japan's concentration in mid-technology sectors is not driven by superior innovation efficiency in these domains. Rather, Japan's R&D investments in mid-tech industries yield fewer patents per dollar spent compared to other countries, while its high-tech R&D is relatively more productive. This reinforces the interpretation that Japan's mid-tech specialization reflects structural factors rather than efficiency-driven optimization.

Figure 16: Relative innovation efficiency (patents per million U.S. dollars of R&D) by country and technology level

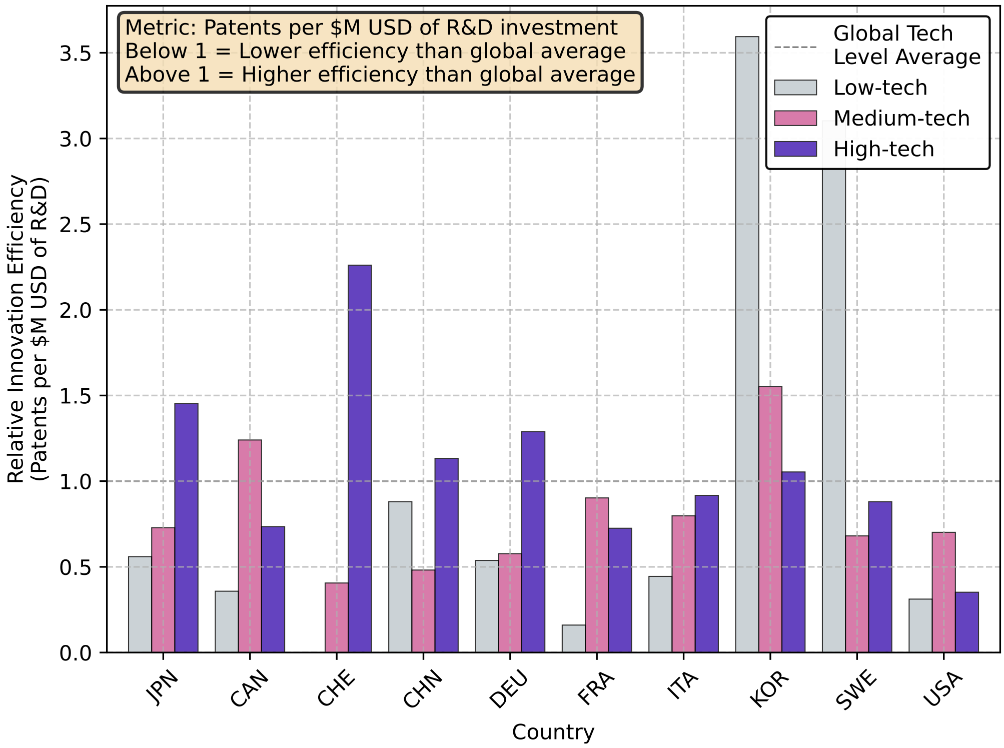

To understand whether public policy can influence the structure of innovation investment, we turn to the composition of government support for research and development. Figure 17, drawn from the (Appelt et al., 2022), highlights a striking feature of Japan's policy framework: the vast majority of R&D tax incentives accrue to large firms, while small and medium-sized enterprises (SMEs) capture only a minor share. This pattern reflects the institutional design of Japan's innovation policy, where tax-based support mechanisms dominate the government's approach to stimulating private R&D.

In Japan, indirect subsidies—mainly in the form of tax credits or deductions—constitute the core of public R&D support (Bahar and Strauss, 2020). These incentives are accessible to all firms that engage in eligible R&D activities, but in practice, they overwhelmingly benefit large, established corporations that already have significant

Figure 17: Distribution of R&D Tax Incentive Support in 2020 by Firm Size

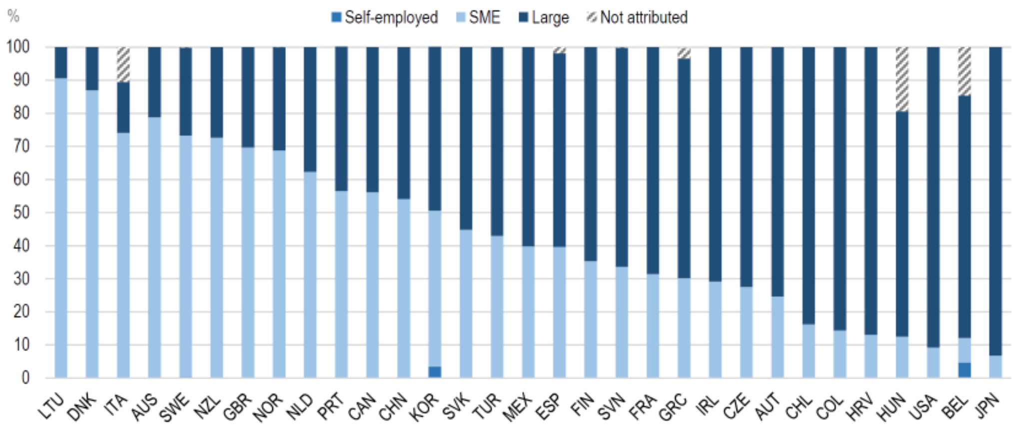

R&D capacity and taxable profits. This design has important implications for innovation dynamism: it reinforces the existing industrial structure rather than encouraging experimentation or entry by smaller, more agile firms. According to OECD estimates, SMEs account for roughly 6 percent of total business R&D.

Yet, small and medium-sized firms—particularly technology-oriented startups and innovative ventures, as distinct from conventional small businesses in traditional sectors—are often the most dynamic actors within innovation ecosystems. They tend to operate closer to the technological frontier in specific niches and are more likely to pursue market-capturing innovation—developing new products or processes not merely to improve efficiency but to create rents through disruption. In environments characterized by fast-paced technological change and competitive startup activity, such firms are key drivers of structural transformation. Their flexibility and appetite for risk allow them to explore ideas that incumbent firms may find too uncertain or misaligned with existing product lines. In Japan's case, the limited participation of SMEs in R&D therefore represents not just an issue of distributional equity, but a lost opportunity for innovation-led growth.

It is important to emphasize that the figure captures only one dimension of public innovation policy—the tax-based (indirect) component. By contrast, other countries, most notably the United States, combine these indirect incentives with substantial direct public funding for R&D through grants, contracts, and mission-oriented programs (Bahar and Strauss, 2020). Agencies such as the National Science Foundation (NSF), the Defense Advanced Research Projects Agency (DARPA), and the Small Business Innovation Research (SBIR) program allocate public funds directly to firms and research organizations, particularly those engaged in high-risk or early-stage innovation. The scale of direct government support varies considerably: approximately 20 percent of U.S. government R&D budgets are allocated directly to enterprises, while in Germany roughly two-thirds of public research institution budgets are funded through government subsidies. This complementary system allows governments such as the U.S. to target frontier technologies and smaller firms more effectively, even if tax incentives themselves exhibit a similar large-firm bias.

In Japan, by contrast, direct funding programs remain comparatively limited in scale, and tax credits remain the primary channel of public R&D support. The result is a policy environment that favors incremental innovation within established corporate ecosystems rather than risk-taking and diversification. Unfortunately, the available data do not allow us to directly observe whether Japan's R&D tax incentives for large firms are being allocated disproportionately to mid-technology rather than high-technology industries. However, given that the majority of tax-based support accrues to large enterprises, we can use complementary evidence from the summary statistics of Japan's Economic Census of 2021 (Statistics Bureau, Ministry of Internal Affairs and Communications, 2021) to understand how industries of different technological intensities are distributed across firm size categories.

Figure 18 presents the composition of establishments and employment by technology level and firm size in Japan for 2021. It is important to note that these data represent the totality of firms in the Japanese economy, not just those engaged in R&D. If we were to restrict the analysis to firms performing R&D, we would expect the shares of mid- and high-technology industries to be significantly higher across both firm-size categories—reflecting both the industrial structure of Japan's R&D-intensive sectors and the design of the incentive system itself. Nonetheless, even without making such assumptions, this figure is highly informative. Among large firms, mid-technology industries account for roughly 23 percent of establishments and nearly 29 percent of total employment, while high-technology industries represent a smaller share (about 15 percent of employment).

These figures imply that, under Japan's current system of broad-based tax incentives, where large firms are the primary beneficiaries, a substantial portion of public R&D support naturally flows toward mid-technology industries, simply because those industries are much more prevalent in the large-firm landscape. In contrast, if a greater share of R&D incentives were directed toward SMEs, the same policy instruments would likely reach a more even mix of mid- and high-technology sectors. Even though high-tech industries constitute a smaller overall share of SME activity, they are proportionally more significant among smaller firms than among large ones, suggesting that shifting the distribution of support toward SMEs could indirectly promote more R&D in high-tech industries.

An additional point is that the scarcity of high-technology firms among SMEs is itself likely endogenous to the current policy environment. Because smaller firms receive very limited public R&D support, their capacity to engage in high-intensity innovation—or even to enter high-tech sectors in the first place—is constrained. Over time, this dynamic

Figure 18: Composition of Establishments and Employment by Technology Level and Firm Size in Japan

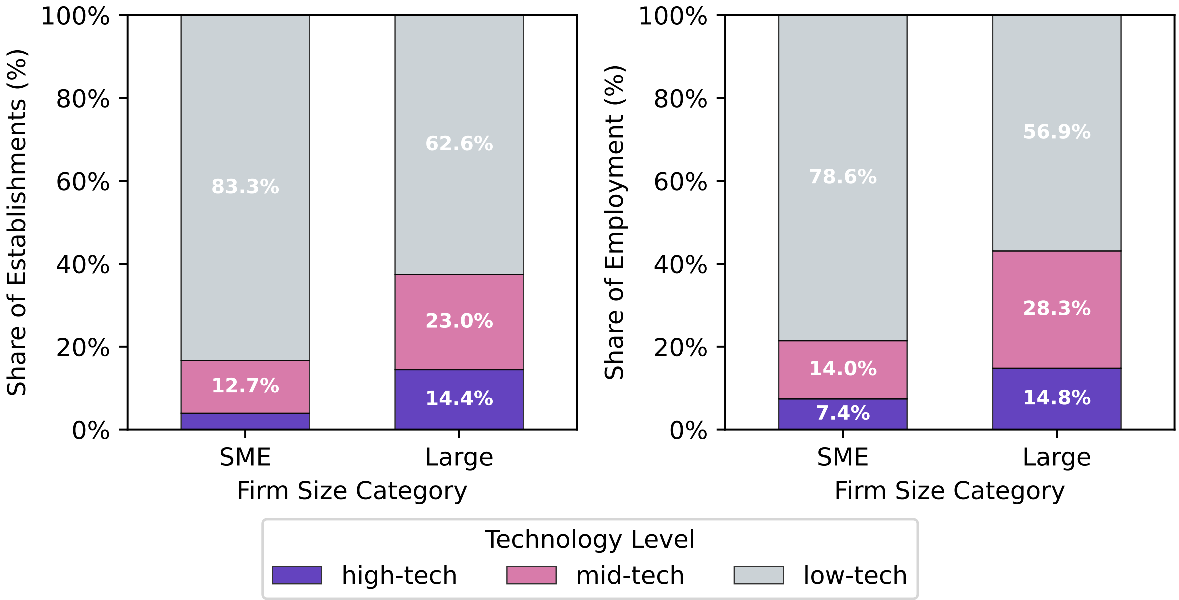

Note: This figure shows how Japanese establishments and employment were distributed across low-, mid-, and high-technology industries in 2021 for both SMEs and large firms. Low-technology industries dominate overall, but among large firms the presence of mid-technology industries is substantial.

reinforces the existing equilibrium: large, established corporations dominate mid-tech R&D activity, while small firms remain peripheral and underfunded in higher-technology domains. In this sense, Japan's firm-size bias in R&D incentives not only reflects but also perpetuates the structural imbalance in the technological composition of its innovation system.

In short, while direct evidence on the technology composition of R&D tax incentives is lacking, the firm-size and industry data together suggest that the structure of Japan's innovation policy mechanically channels public support toward mid-technology activities. Rebalancing incentives to expand SME participation in R&D could, therefore, have the dual effect of improving inclusiveness and shifting innovation toward higher-technology sectors with greater productivity potential.

Policy Window: Japan's 17–Sector Strategy

In this context, the Japanese government's recently announced strategy prioritizing 17 strategic innovation domains—including semiconductors, AI systems, biotechnology, advanced energy technologies, cybersecurity, aerospace, and quantum computing—represents a potentially important shift in Japan's innovation policy[12]. On paper, this initiative is

[12]For a press report on the announcement, see https://www.asahi.com/ajw/articles/16137063.

precisely the kind of directional reallocation that our analysis suggests is necessary. It places explicit emphasis on frontier technologies that generate large productivity spillovers and signals a willingness to move beyond incremental improvements in mature manufacturing industries.

To gauge the potential of this strategy, it is instructive to compare Japan's current R&D intensity across these sectors with that of the United States. Figure 19 plots R&D intensity (R&D expenditures as a share of GVA) for the sectors most closely related to the government's 17 strategic domains, highlighting their technology classification (low, mid, high). As observed, high-technology sectors such as pharmaceuticals (C21) and information services (J62, J63) exhibit substantially higher R&D intensity in the United States than in Japan. Conversely, Japan's relative strengths remain concentrated in mid-technology industries where R&D intensity is closer to U.S. levels but aggregate productivity growth is weaker.

Figure 19: Average R&D Intensity in Strategic Sectors: Japan vs. United States

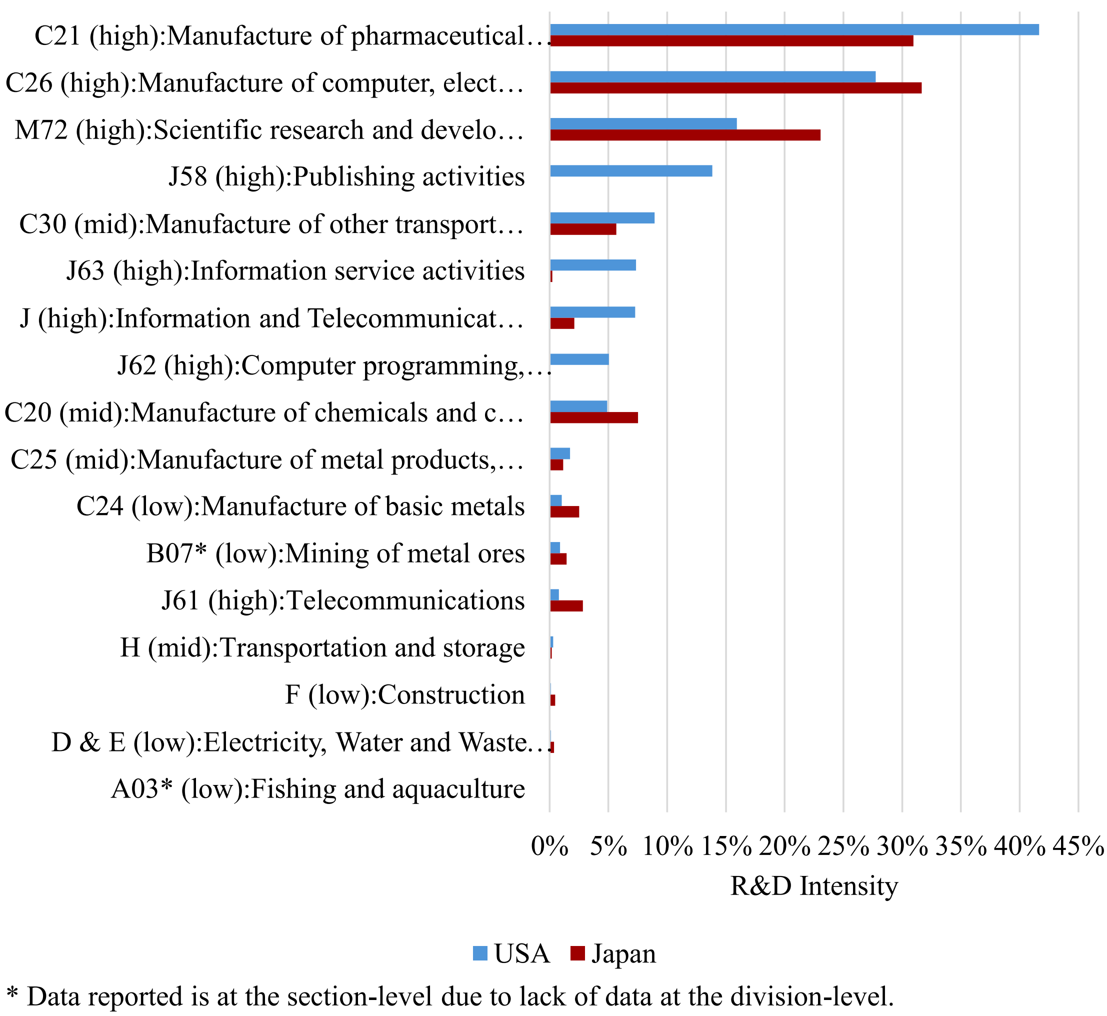

Note: Sectors are shown at the section level due to data limitations. Technology classifications (low, mid, high) follow the taxonomy described earlier. Data source: Author calculations using OECD STAN.

This pattern reinforces a central message of the paper: Japan's R&D system remains oriented toward technologically sophisticated but structurally mature industries, while the United States invests far more heavily in frontier domains closely aligned with the 17–sector strategy. If the new initiative successfully shifts resources toward these high-technology sectors—where Japan currently underinvests relative to its peers—it could be a decisive step toward breaking out of the middle-technology trap.

The success of the 17–sector strategy, however, depends critically on implementation. If the program is executed primarily through the existing tax-based R&D incentive system, large incumbent firms will capture most of the benefits, and Japan's technological allocation may remain largely unchanged. For the initiative to fulfill its potential, support must be explicitly targeted to high-technology frontier activities rather than distributed uniformly across designated sectors, and a substantial share of funding must be accessible to SMEs, startups, and early-stage technology firms rather than being captured disproportionately by large incumbents.

Viewed through the lens of Figure 19, these conditions take on added significance. The strategic sectors where Japan lags the most are precisely those where SMEs and young firms historically play a central role in generating breakthroughs. Without mechanisms that enable these actors to participate meaningfully in the 17–sector strategy, the initiative risks becoming another well-intentioned policy that reinforces existing patterns: high effort, limited reallocation, and modest gains in aggregate productivity. If implemented with genuine reallocation in mind, the proposed government strategy could serve as an institutional vehicle for helping Japan escape the middle-technology trap. If not, it may simply broaden the scope of existing incentives without altering the underlying distribution of innovation effort.

Concluding Remarks and Policy Insights

Japan remains one of the most technologically capable countries in the world, yet its labor productivity has barely grown in over two decades. The analysis in this paper suggests that this is not due to insufficient innovation effort—R&D intensity is among the highest in the OECD—but rather to the *allocation* of that effort. Japanese firms invest heavily in innovation, but disproportionately in mid-technology industries that are technologically sophisticated yet structurally mature. These sectors continue to produce incremental advances, but they no longer generate the transformative productivity effects associated with frontier technological growth.

Our empirical results reveal three central findings. First, more than half of Japan's business R&D is allocated to mid-technology industries, a concentration higher than in any comparable economy. Second, high-technology sectors globally generate the largest contributions to productivity growth, yet Japan devotes a comparatively small share of R&D to these frontier domains. Third, the marginal productivity return to R&D in Japanese mid-tech industries appears weaker than the equivalent return in other advanced economies, suggesting potential misallocation. Together, these facts point to the presence of a middle-technology trap: Japan is too advanced to rely on catch-up growth, yet too concentrated in mature industries to generate disruptive, frontier-expanding innovation.

A structural mechanism contributes to this equilibrium. Japan relies overwhelmingly on indirect, tax-based R&D incentives, which largely benefit large incumbent firms. Because these incumbents are concentrated in mid-tech industries, public support tends to reinforce the status quo rather than diversify innovation into emerging technological domains. Small and medium-sized enterprises—where high-tech formation is more likely to originate, based on the experience of other countries—receive only a small share of total R&D support. This is not simply a distributional imbalance; it is a constraint on frontier innovation. The absence of a broad high-tech SME base therefore limits Japan's innovation option set.

The policy implications of our analysis follow directly from this diagnosis. First, shifting the marginal yen of R&D toward frontier sectors—including semiconductors, AI hardware and software, biotechnology, advanced materials, quantum technologies, and green energy—is likely to raise productivity where spillovers are largest. Second, expanding support for high-tech SMEs and mid-risk innovation through instruments such as refundable R&D credits, competitive grants (*i.e.*, SBIR-style grants), procurement pathways for emerging technologies, and regulatory simplification would broaden the set of actors able to explore new technological opportunities. Third, increasing the relative weight of direct funding—mission-oriented programs and targeted grants—vis-à-vis tax incentives would allow the state to underwrite riskier, exploratory projects that incumbent firms might otherwise underinvest in. International experience suggests that dedicated agencies can play a central role in this regard: institutions such as the U.S. DARPA or Japan's New Energy and Industrial Technology Development Organization (NEDO) demonstrate how publicly funded, high-risk research programs can absorb technological uncertainty, coordinate long-horizon projects, and accelerate the emergence of new sectors. Finally, the newly announced 17–sector strategy can serve as a concrete reallocation mechanism rather than a simple subsidy amplifier: if its instruments are designed to channel resources toward frontier activities and to be accessible to younger, smaller firms, it can help redirect Japan's innovation system away from its current mid-technology bias.

Crucially, expanding R&D support for SMEs and high-tech entrants does not weaken Japan's major corporations; rather, it can strengthen them. A more dynamic SME layer creates exploratory technologies, intellectual property, and proof-of-concept products that large firms can later license, acquire, or scale. Innovation ecosystems are pipelines: new technology is explored at the edge and industrialized at the core. A policy regime that expands the exploratory frontier therefore raises the productivity potential of incumbent firms instead of competing with them.

Japan's productivity challenge is thus not a paradox but a misallocation problem. The country does not need to innovate more; it needs to innovate *differently*. By broadening the innovation base, empowering smaller firms, and directing additional R&D toward frontier technologies, Japan can convert its technological sophistication into productivity growth once again. The 17–sector strategy creates a window for doing so. Whether that window becomes transformative depends on how strongly it is used to reshape the allocation of innovation effort.

References

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2):323–351.

Aghion, P. and Howitt, P. (2009). *The Economics of Growth*. The MIT Press, Cambridge, MA.

Appelt, S., Günter, L., and Galindo-Rueda, F. (2022). Oecd r&d tax incentives database report, 2022 edition. OECD Directorate for Science, Technology and Innovation.

Arora, A., Belenzon, S., Cioaca, L., Sheer, L., and Shvadron, D. (2024). Discern 2: Duke innovation & scientific enterprises research network.

Bahar, D., Arcay, G., Daboin Pacheco, J., and Hausmann, R. (2024). Japan's economic puzzle. Working Paper CID Working Paper No. 442, Harvard Growth Lab, Center for International Development, Harvard University. Revised July 2024.

Bahar, D. and Strauss, S. (2020). Innovation and the transatlantic productivity slowdown: A comparative analysis of r&d and patenting trends in japan, germany and the united states. Global working paper, The Brookings Institution, Washington, DC. Brookings Global Working Paper.

Chacua, C. (2019). patstat2018b_load. https://github.com/cchacua/patstat2018b_load. Version 20190217. Computer software.

Chacua, C. (2023). green_patstat. https://github.com/cchacua/green_patstat. Version 20230818. Data set.

Fernández-Villaverde, J., Ventura, G., and Yao, W. (2023). The wealth of working nations. NBER Working Paper 31914, National Bureau of Economic Research, Cambridge, MA. Revised August 2024.

Fuest, C., Gros, D., Mengel, P.-L., Presidente, G., and Tirole, J. (2024). Eu innovation policy: How to escape the middle technology trap. Technical report, CESifo, BU Institute for European Policymaking, Bocconi University, and Toulouse School of Economics, Munich, Milan, Toulouse. Economics for the Common Good, Policy Report No. 50.

Hausmann, R., Hidalgo, C. A., Bustos, S., Coscia, M., Simoes, A., and Yıldırım, M. A. (2014). *The Atlas of Economic Complexity: Mapping Paths to Prosperity*. The MIT Press.

Miguélez, E., Raffo, J., Chacua, C., Coda-Zabetta, M., Yin, D., Lissoni, F., and Tarasconi, G. (2019). Tied in: The global network of local innovation. *WIPO Economic Research Working Paper*, (58).

Statistics Bureau, Ministry of Internal Affairs and Communications (2021). Economic census for business activity 2021: Summary statistics (japan). e-Stat: Portal Site of Official Statistics of Japan. Dataset retrieved from https://www.e-stat.go.jp/en.

Appendix A: Technology Classification Methodology

This appendix provides details on the methodology used to classify economic sectors by technology intensity level (high-tech, mid-tech, and low-tech) according to the ISIC Rev. 4 industry classification system. The final consensus classification presented in the main text was constructed by comparing and synthesizing multiple classification schemes.

We created these classifications using data from several sources and applying different criteria to categorize industries. These sources capture distinct dimensions of technological intensity: some focus on the volume and growth dynamics of patent activity within technology classes, while others measure the recency and age structure of innovation. Additional methods incorporate economic complexity metrics. We also include LLM-based assessments that evaluate industries from multiple perspectives—considering both the technological sophistication of production inputs (such as workforce skills and advanced processes) and the innovation embodied in outputs (such as product complexity and novelty). By leveraging these diverse methodologies, we aimed to create a robust classification that captures the multifaceted nature of technological intensity across sectors.

To reconcile differences across classification sources and establish a single consensus classification, we implemented a majority voting mechanism. For each industry, we aggregated the classifications from all available sources and selected the technology level that received the most votes. In cases where no clear majority emerged (*i.e.*, tied votes), we applied a progressive tie-breaking rule that prioritizes higher technology classifications when ambiguity exists.

The voting process generates a confidence score for each industry based on the proportion of sources agreeing with the final classification. Industries with high agreement across sources (more than or equal to 80% consensus) provide strong validation of their technology intensity, while those with lower agreement may reflect genuine ambiguity in their technological characteristics or differences in how various methodologies weight technological factors.

Figure 20 presents a heatmap comparing technology level assignments across all classification methods for industries included in our balanced panel dataset. Each row represents an ISIC industry code, and each column corresponds to a different classification methodology. The color coding indicates the assigned technology level: high-tech (dark purple), mid-tech (pink), and low-tech (light gray).

The heatmap reveals several patterns. First, there is substantial agreement across methods for industries at the technological extremes—those clearly high-tech (*e.g.*, pharmaceuticals, computer manufacturing) or definitively low-tech (*e.g.*., basic textiles, food

Figure 20: Technology Classification Comparison Across Methods

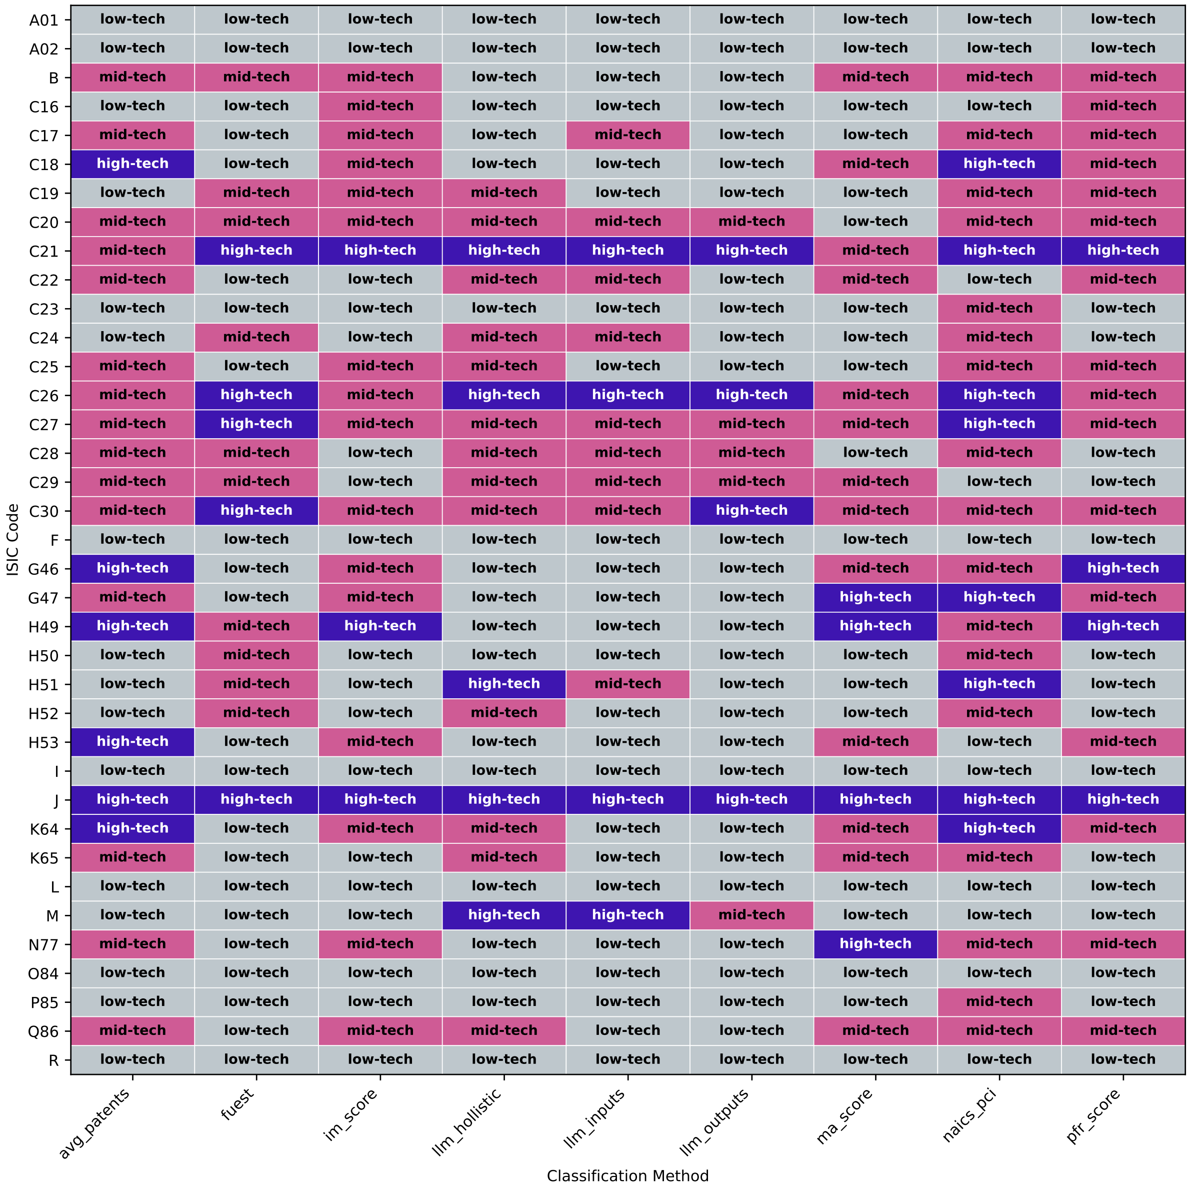

Note: This figure presents a heatmap comparing technology level assignments across all classification methods for industries included in our balanced panel dataset. Each row represents an ISIC industry code, and each column corresponds to a different classification methodology.

processing). Second, the greatest divergence appears in mid-tech industries, where different methodologies may emphasize different aspects of technological intensity. Third, some industries show splits between adjacent categories (*e.g.*, high-tech vs. mid-tech), while cross-category disagreements (high-tech vs. low-tech) are rare, suggesting that classification methods capture a consistent underlying technological gradient.

To quantify the consistency across classification schemes, we computed pairwise agreement rates between all methods. Figure 21 displays these agreement rates in a symmetric matrix, where each cell indicates the proportion of industries receiving identical classifications from two methods.

Figure 21: Pairwise Agreement Rates Between Classification Methods

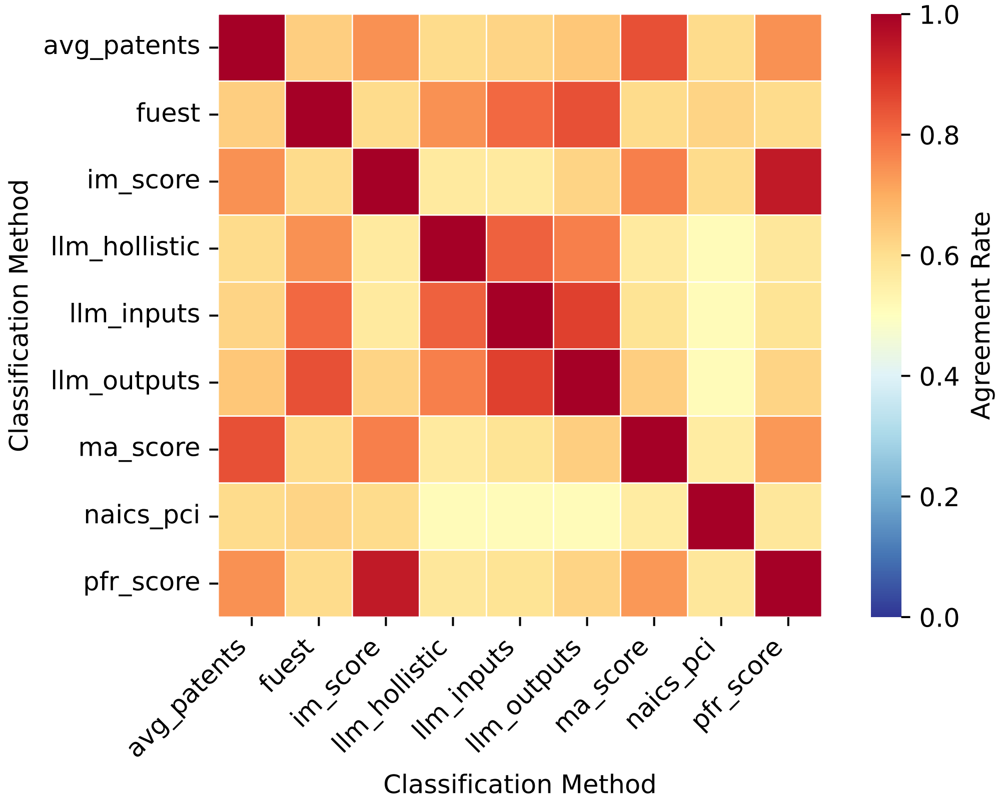

Note: This figure presents a pairwise correlation matrix across all classification methods. Each row and column represents a classification methodology.

The agreement matrix shows moderate to high concordance across most classification pairs, with average pairwise agreement of approximately 60-70%. This pattern provides confidence that our consensus classification captures a genuine technological dimension that is robust to methodological choices.

Appendix B: Regression Results

This appendix presents the main regression analysis examining the relationship between technology levels and productivity growth components. We test whether our results are sensitive to the choice of technology classification scheme by running the same regression specification across nine different methods of classifying industries as high-tech, mid-tech, or low-tech.

For each classification scheme, we estimate the following model:

$$Y_{i,j} = \beta_0 + \beta_1 \text{HighTech}_i + \beta_2 \text{MidTech}_i + \beta_3 \text{GVAShare}_{i,j} + \gamma_j + \varepsilon_{i,j} \tag{2}$$

The variable $Y_{i,j}$ denotes the contribution of sector i to the annual average growth of labor productivity in nation j. The variables HighTech_{ij} and MidTech_{ij} function as binary indicators to determine whether the industry is categorized within each technological group, as per the classifications previously delineated. The variable $\text{GVAShare}_{i,j}$ serves as a control variable, representing the proportionate significance of industry i within the economy. Furthermore, the parameter γ_j accounts for country-specific fixed effects, thus controlling for macroeconomic and institutional disparities across different countries. The dataset employed in this analysis encompasses the same nations and industries detailed in the data section. Notably, to evaluate long-term effects, we compute the contribution of industry i to productivity growth in country j across the entire period from 2000 to 2021. Consequently, the regression analysis utilizes 885 observations.

The table below presents results using total contribution to productivity growth as dependent variable with columns representing different classification schemes. Consistent signs and significance levels across classification schemes provide evidence that the relationship between technology level and productivity growth is robust to how we classify industries.

Table 2: Regression Results: Total Contribution (Robustness Across Classifications)

	Dependent variable: total_contribution								
	Avg. Patents	Fuest	IM Score	LLM Holistic	LLM Inputs	LLM Outputs	MA Score	NAICS-PCI	PFR Score
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constant	0.028*	0.037**	0.027*	0.017	0.026*	0.018	0.022	0.037**	0.030*
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
High-Tech	0.014	-0.011	0.024**	0.037***	0.046***	0.011	0.025**	-0.014	0.035***
	(0.009)	(0.010)	(0.012)	(0.010)	(0.010)	(0.010)	(0.011)	(0.009)	(0.010)
Mid-Tech	-0.019***	-0.036***	-0.015**	-0.002	-0.028***	0.017*	-0.005	-0.023***	-0.023***
	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.009)	(0.007)	(0.007)	(0.007)
Industry Share	0.479***	0.395***	0.499***	0.542***	0.442***	0.549***	0.474***	0.465***	0.445***
	(0.087)	(0.091)	(0.087)	(0.089)	(0.088)	(0.087)	(0.088)	(0.088)	(0.086)
Technology Level	Dummies	Dummies	Dummies	Dummies	Dummies	Dummies	Dummies	Dummies	Dummies
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control: Ind. Share	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	885	885	885	885	885	885	885	885	885
R^2	0.261	0.267	0.259	0.263	0.284	0.252	0.255	0.257	0.276
Adjusted R^2	0.239	0.245	0.237	0.241	0.262	0.230	0.233	0.234	0.254
Residual Std. Error	0.093	0.093	0.093	0.093	0.092	0.094	0.094	0.094	0.092
F Statistic	11.670***	12.048***	11.549***	11.804***	13.078***	11.136***	11.307***	11.415***	12.589***

Note: *p< 0.1; **p< 0.05; ***p< 0.01.